

# ANGSTY METAETHICS

Joshua Blanchard

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Philosophy in the College of Arts and Sciences.

Chapel Hill  
2018

Approved by:

Geoffrey Sayre-McCord

Susan Wolf

Ram Neta

Russ Shafer-Landau

Ryan Preston-Roedder

© 2018  
Joshua Blanchard  
ALL RIGHTS RESERVED

## ABSTRACT

Joshua Blanchard: Angsty Metaethics  
(Under the direction of Geoffrey Sayre-McCord)

Some disagreements evoke *philosophical angst*: the judgment that the truth of some thesis is essential for the meaning or intelligibility of our lives, combined with the worry that it might be false. Chapter 1 develops an account of this attitude, its phenomenology, and its normative status in contrast to related attitudes. Philosophical angst illuminates what is at stake in debates ranging from the metaphysics of reasons to the relationship between God and meaning.

Chapter 2 defends the motivation and coherence of the project. I give a general argument for expecting that there are significant evaluative differences in the implications of competing metaethical theories: radically different accounts of important domains probably differ in evaluative upshot. I then respond to three puzzles for making evaluative judgments about metaethics.

Chapter 3 defends *pro-realism*, the view that it is much better if moral realism is true rather than any of its rivals. First, moral realism vindicates the dignified moral status described by the best normative moral theories, and so it is much better if realism rather than nihilism is true. Second, moral realism secures a desirable *independence* for moral justification that is different in kind from anti-realistic independence.

Chapter 4 rebuts arguments that it is better if antirealism is true. First, there are reasons for thinking that morality itself would be worse if realism were true. Second, there are moral reasons for thinking we shouldn't endorse realism. I argue that such arguments either rely on implausible grounding principles or they overgeneralize to all metaethical projects.

Although I have raised questions of angst and value, one might think that this is irrelevant to truth. Chapter 5 shows that this reasonable impulse is mistaken. Inferences from *It is better if p* to *It is the case that p* are appropriate in contexts of inquiry in which theorists are responsible for accommodating not only the non-evaluative features of some domain, but the evaluative features as well. Metaethics is one such context.

## ACKNOWLEDGEMENTS

For helpful discussion and feedback on both specific portions and the central themes of this dissertation I thank Benjamin Bagley, Lindsay Brainard, Spencer Case, Terence Cuneo, Raff Donelson, Luke Elson, David Faraci, Krasi Filcheva, Eli Hirsch, Zoë King, Barry Maguire, Ram Neta, Derek Parfit, Ryan Preston-Roedder, Russ Shafer-Landau, Keshav Singh, Larisa Svirsky, Silvan Wittwer, Susan Wolf, Monique Wonderly, and Alex Worsnip. I would also like to acknowledge the audiences at the 2017 Rocky Mountain Ethics Congress, the 2018 Great Lakes Philosophy Conference, and my colleagues in UNC's Work in Progress Series and Dissertation Research Seminar.

I am especially grateful to my mentor and dissertation chair, Geoffrey Sayre-McCord, whose dedication to conceptual clarity, intellectual charity, and good coffee will always be a model for me.

I am also grateful to my parents, Scott and Judy Blanchard, who fostered a spirit of open but challenging philosophical dialogue in the home and have always supported my academic pursuits.

Finally, I thank Bethany Blanchard, whose love and diligence afforded me the privilege to spend time being angsty about arcane matters in moral philosophy.

## TABLE OF CONTENTS

CHAPTER 1: PHILOSOPHICAL ANGST.....	1
Philosophical Angst: Its Nature.....	1
Philosophical Angst: Its Normative Status.....	6
Insiders and Outsiders.....	17
Philosophical Angst and Contemporary Philosophy.....	21
Conclusion.....	27
CHAPTER 2: METAETHICAL VALUE JUDGMENTS.....	28
Metaethics Matters: Some Negative Arguments.....	29
Metaethics Matters: A Positive Argument.....	32
Some Puzzles About Metaethical Value Judgments.....	36
Conclusion.....	44
CHAPTER 3: MORAL REALISM AND PHILOSOPHICAL ANGST.....	46
Philosophical Angst.....	47
How to Be Angry About Moral Realism.....	50
Conclusion.....	76
CHAPTER 4: PREFERRING MORAL ANTIREALISM.....	77
Anti-Moral Realism.....	78

Responding to Anti-Moral Realism.....	87
Are Normative and Metaethics Normatively Unrelated?.....	95
Conclusion.....	96
CHAPTER 5: VALUE AS A GUIDE TO TRUTH IN METAETHICS.....	98
Traditional Objections Rejected.....	99
Some Exotic Value-Truth Linking Premises.....	102
Value as a Guide to Truth in Metaethics.....	104
Comparative Advantages of the View.....	109
Conclusion.....	118
REFERENCES.....	119

# Chapter 1: Philosophical Angst

This chapter is about a particular way in which philosophical claims sometimes matter. Under the right conditions, such a claim can be the appropriate object of a distinctive concern, what I call *philosophical angst*. Through a series of examples and comparisons to ordinary phenomena like grief, I show that the concept captures something both familiar and important that is also uniquely philosophical. Achieving a better understanding of philosophical angst helps us to see why philosophy itself matters, and it highlights underexplored contours of otherwise longstanding debates.

## §1 Philosophical angst: its nature

Angsty philosophers work on topics as diverse as the existence of God, the immortality of the soul<sup>1</sup>, the freedom of the will, the nature of the mind<sup>2</sup>, and even the metaphysics of modality.<sup>3</sup> In each of these domains, philosophers disagree about what is true, but they sometimes also disagree about the evaluative significance of competing possibilities—that is, whether and why it matters what is true. Such concerns are not, of course, limited to professional philosophers. We

---

<sup>1</sup> Eli Hirsch (2009, 2, 20) says that both death and radical deception are “ineffably horrible” and “undermine everything that I have ever cared about in my life.”

<sup>2</sup> Jerry Fodor (1989, 77) says, winsomely, that “if it isn’t literally true that my wanting is causally responsible for my reaching, and my itching is causally responsible for my scratching, and my believing is causally responsible for my saying.... If none of that is literally true, then practically everything I believe about anything is false and it’s the end of the world.” Thanks to Matthew Kotzen for this reference.

<sup>3</sup> Robert Adams (1979) argues that modal realism would lead to total “moral indifference.”

often know philosophical angst when we see and feel it, but the questions of its nature and appropriateness have not received sustained attention. At the very least, I want to understand and show how angst can be appropriate and rationally intelligible for a person, but I will also explore whether a case can ever be made that it is incumbent upon everyone to be angsty. It is one thing to show that the angsty among us have intelligible attitudes, but quite another to show that the angst-free are somehow deficient.

What do I mean by “philosophical angst”? The word “angst” is familiar both as a term of art in existentialist or continental philosophy (where it sometimes goes by “anxiety” or “dread”) and as a descriptor for some related feelings stereotypically attributed to young adults. Some general remarks and examples will help to clarify the particular phenomenon that I have in mind.

In a well-known parable from *The Gay Science*, Friedrich Nietzsche offers what I regard as a paradigmatic example of philosophical angst.<sup>4</sup> In the parable, “many of those who did not believe in God were standing around together” in a marketplace.<sup>5</sup> A “madman” runs through the crowd screaming, astonished that no one shares his dread toward the prospect of atheism. They (and here the madman includes himself) have “killed” God; they have abandoned any apparent belief in and commitment to God or the idea of God.<sup>6</sup> Despite taking on atheism as a philosophy, they have not fully appreciated what things are really like if atheism is true. The madman’s problem with the crowd is not that they are atheists *per se*. His problem is that they do not appreciate the implications of their atheism. In a hopeless effort to jolt them out of apathy, the

---

<sup>4</sup> Nietzsche (2001, 119-120).

<sup>5</sup> Ibid., 119.

<sup>6</sup> See Nietzsche’s own interpretation in *ibid.*, “The greatest recent event – that ‘God is dead’; that the belief in the Christian God has become unbelievable – is already starting to cast its first shadow over Europe” (343).

madman bombastically berates the crowd. He declares that the “holiest and the mightiest thing the world has ever possessed has bled to death under our knives.”<sup>7</sup> Existing without God is like “drink[ing] up the sea,” or straying “as though through an infinite nothing,” and there is no obvious way for us to “console” or “clean” ourselves in light of this calamitous fact.

The madman seems to think that it is horrible (or, at the very least, colossally unsettling) that God does not exist, and further that someone attuned to this fact should be disturbed by the prospect of atheism. He asks rhetorical questions about the seemingly Sisyphean task of replacing God—for example, “Must we ourselves not become gods simply to appear worthy of it?”—and these illustrate what I take to be the core cognitive judgment constitutive of philosophical angst: It is very important that a philosophical claim be true, because there is no existentially adequate replacement for what the claim posits. What I mean by “existential” in this context has nothing to do with quantification. Rather, something is existentially important when it is important to the overall meaning and intelligibility of our lives. So, to be philosophically angsty is to think that what some philosophical claim posits—most richly, the way that some philosophical perspective pictures the world—is irreplaceable with respect to its contribution to the meaning or intelligibility of our lives, but that it is or might be false. Notice that, like Nietzsche’s madman, one need not believe the claim or perspective in question. One must simply regard its falsity (whether actual or merely possible) with horror.

Irreplaceability, or more specifically what I will call “existential irreplaceability,” is a necessary but not a sufficient condition for philosophical angst. To illustrate, consider contemporary debates about the existence and value of God. Some philosophers, for example, Mark Johnston, Ronald Dworkin, and other religious naturalists, think that God’s existence

---

<sup>7</sup> Ibid., 120.

would play an existentially *significant* role in human life, but that we have adequate replacements; in the case of Johnston, the replacements are humanity's highest ideals; in the case of Dworkin, the replacement is the cosmos writ large as an object of the traditionally religious attitudes.<sup>8</sup> In order to be properly angsty about the non-existence of God, one must not only think that God's existence would satisfy deep human needs (Johnston and Dworkin agree that it would), but that only God's existence could do so (here, they demur). If humanity's highest ideals do just as well (or well enough), then someone who recognizes this fact cannot be intelligibly angsty about the falsity of theism.

Likewise, the fact that something plays an existentially significant role does not mean that it or the role it plays is important in a sense that entails its being good; for this reason, mere existential irreplaceability is not a sufficient condition for philosophical angst. Perhaps, as Nietzsche himself ultimately thinks, the role that God has played in human life and thought is, in the end, bad or ignoble in some way. Accordingly, someone who thinks that God uniquely plays an existentially significant role in human life may not yet be angsty about the falsity of theism. They must also think that God plays an important positive role, and that it is bad if nothing plays that role. Call this more value-laden sense of irreplaceability, "existentially important irreplaceability."

A complex judgment of existentially important irreplaceability, though necessary, is still not a sufficient condition for philosophical angst. Such a judgment, as so far described, involves two cognitive components. First, one judges that something plays an irreplaceable (not just significant) role in human life. Second, one judges that what is irreplaceable is also very important in a positive sense. But nothing has been said so far about the phenomenology

---

<sup>8</sup> Johnston (2009); Dworkin (2013).

naturally attendant to these judgments. For all I have said, someone could make a judgment of existentially important irreplaceability and simply move on with their lives.<sup>9</sup> While it is tempting to suggest a necessary link between judgments of existentially important irreplaceability and felt anxiety, the temptation should be resisted. Felt anxiety is only appropriate when a third condition is met: one comes to think that what is existentially important and irreplaceable either does not or might not exist. Here, there may indeed be a necessary normative linkage--those who judge that something existentially important and irreplaceable does not or might not exist plausibly ought to in turn feel anxiety, or are at least rationally licensed to do so.

It is worth independently emphasizing that, like a judgment of existentially important irreplaceability, felt anxiety itself is not a sufficient condition for philosophical angst. Someone might feel anxiety toward the potential loss or absence of something that they do not judge to have existentially important irreplaceability. For example, an otherwise angsty atheist may believe that there are naturalistic ways to satisfy religious desires and sentiments, but stubbornly or otherwise refuse to pursue them--wallowing in the absence of God despite the fact that, by their own lights, the wallowing lacks cognitive justification. I do not commit here to thinking that such a person is being irrational or unintelligible. For now, I merely stipulate that they lack philosophical angst.

In sum, philosophical angst has both distinctly cognitive and conative components. The cognitive component is a complex judgment of existential irreplaceability about something important that one thinks does not or might not exist--whether it is God, moral truth, freedom,

---

<sup>9</sup> As Susan Wolf (2010, 29 n.14) comments in her work on meaning: "Many people are upset by the thought that they are mere specks in a vast universe. They are upset, that is, by their smallness, their inability to make a big and lasting splash. My remarks--aimed at reminding them of the quality, not the quantity, of their contribution to the universe--do not speak directly to this concern. Such people will just have to get over it. Their desire is unsatisfiable."

or anything else. The conative component is a kind of anxiety that is naturally attendant to or licensed by this judgment. Putting all of this together, then, to be philosophically angsty is both to judge that something is existentially important and irreplaceable, and to worry about its non-existence.

## **§2 Philosophical angst: its normative status**

Now that my account of philosophical angst is on the table, we can ask the following normative questions. Under what conditions is philosophical angst intelligible? And is it ever rationally or otherwise obligatory? Here is an easy answer: to the extent that someone sees, or is in a position to see, that something is both existentially important and irreplaceable, then it is appropriate for that person to regard its possible or actual non-existence with philosophical angst. This answer is much *too* easy. For one thing, it overgeneralizes. There are plenty of cases where something's existentially important irreplaceability is relativized to certain individuals or groups, as well as cases where it seems appropriate to move beyond one's angst after some finite period of time—for example, in cases of grief, discussed below. In such cases, a person might see that something is both existentially important and irreplaceable, but be outside of the group (or timeframe) in which philosophical angst is appropriate. Yet, what makes philosophical angst distinctly *philosophical* is not just that its object is stereotypically so, but that its normative purport is universal. So, while we need the account to preclude false positives, we also need it to yield the result that cases of appropriate philosophical angst purport to bear, in some sense, on everyone who is properly attuned to the relevant details. A series of examples, drawn from both existing literature and imagination, illustrates the need for some additional taxonomy that militates against the simple answer. My methodology is to use overlapping but distinct

phenomena to better circumscribe the nature, and especially the normative status of, philosophical angst.

### §2.1 *Grief*

Grief over the death of a loved one bears some salient similarities to philosophical angst. In paradigmatic experiences of grief, one judges that the contribution that the lost person makes to one's life is good, that the person is irreplaceable, and even if one does not feel anxiety *per se*, the loss is accompanied by a characteristic conative component—a kind of sorrow, perhaps attendant to felt emptiness. Such felt emptiness may resemble much more general feelings that the world is without value, or that there is no point in living. I doubt that there is a single phenomenology of grief; in any case, the preceding one-sentence gloss is hardly adequate. Indeed, how grief is experienced and expressed in practices of mourning likely varies between both individuals and cultures.<sup>10</sup> But what I want to focus on here is grief's normative status.

The appropriateness of grief is relative in the sense that it depends in part on the relationship that the bereaved bears to the one who is lost. When a child's beloved parent dies, for example, the child's grief is appropriate in a way that a stranger's grief would not be. Though it may be appropriate for a stranger to feel and express sympathy and sorrow upon hearing of the parent's death (and the child's loss), it is not typically appropriate for the stranger to *grieve* the death. There are many possible explanations for this asymmetry, but the one that I consider here is that, for the stranger, the one who is lost did not play a particular and important role in their life, and certainly did not *irreplaceably* play such a role. Strangers, *qua* strangers, are unimportant and replaceable in ways that beloved parents are not. (I do not, of course, mean that they are

---

<sup>10</sup> See, for example, Amy Olberding (1997).

unimportant and replaceable full stop, morally speaking.) The upshot is that the particularity of the deceased is inextricably linked to who and what they were to those who grieve.

The terrain of special relationships *vis-à-vis* grief is vast and complex: we may grieve the loss of neighbors, lovers, teammates, idols, friends, teachers, rivals, children, leaders, and myriad others; grief may exhibit an idiosyncratic felt character in each case. But no matter the felt character, the appropriateness of grief is relationship-dependent. It is not only *not* appropriate for me to grieve a stranger's child in the way that I would grieve my own, it is positively *inappropriate*. And except in extraordinary cases (if any), it is positively inappropriate for a parent *not* to grieve their own child.

So, grief is an attitude constituted by judgments of the loss of something important and irreplaceable, with an attendant negative felt character, but is such that its causes do not make it incumbent upon everyone. Indeed, it is not only not incumbent upon everyone to have it, but it is incumbent upon everyone not to have it unless they exist in the right kind of relationship.

Grief over personal loss is not an instance of philosophical angst, first and foremost because (with the very interesting exception of a personalistic theism to which I return, below) the loss of a person is not sufficiently general in character. The kind of loss in question is itself qualitatively different. One loses a person who really lived, but one does not lose (or worry about the absence of) the truth of a philosophical claim in this way.<sup>11</sup> The loss that motivates philosophical angst is closer to the loss of a belief or, more richly, the loss of a world picture. It is not as if some angst-inducing thing or idea goes from true to false or *vice versa*, in the way that a

---

<sup>11</sup> Recall that I am discussing grief over the death of a loved one. There is an experience recognizable as a form of grief involved in coming to realize that a close relationship has come to an end (as when a friendship or marriage comes to an end), or that an apparent relationship was an illusion (as when one discovers longstanding and unrepentant betrayal by one's spouse). I regard these kinds of grief as derivative or metaphorically extended from the core notion of grief in response to death.

person moves from life to death. Rather, one either moves from *thinking* that something is true, to doubting its truth, or (in perhaps the most angst-ridden state) to thinking that it is false.

The characteristic generality of philosophical attitudes yields a second difference between grief and paradigmatic instances of philosophical angst, namely that grief is relationship-dependent in a way that philosophical angst is not. While appropriate grief does respond to something objectively bad (the death of a person), the response is licensed only in certain relational (what we might call “subjective”) contexts.

Despite not being a form of philosophical angst, grief nevertheless illuminates an interesting possibility with respect to philosophical angst. There is a kind of attachment constitutive of many losses that appropriately lead to grief that arguably has analogues in the philosophical case. Monique Wonderly offers the best account of this phenomenon, what she calls “security-based attachment.” According to Wonderly, we are attached in this way when the object of our attachment is some “non-substitutable particular” that we experience as a “felt need,” such that its presence is a necessary and sufficient condition for a certain feeling of safety or at-home-ness in the world.<sup>12</sup> Wonderly’s account is tailored to losses that are less general than putative philosophical truths. Indeed, it is awkward to describe a philosophical object, or what some philosophical claim posits (for example, moral value, God, or a kind of freedom), as a “non-substitutable particular,” albeit not wholly off the mark. If the philosophical claim posits something of a *sui generis* nature, then there is at least a sense of particularity enjoyed by the claim, though a sense quite different from the way in which one’s mother is a non-substitutable particular. (“Freedom of the will is irreplaceable” has a very different ring from “Your mother is irreplaceable.”) Nevertheless, the analogy between the “felt need” of a person and the felt need of

---

<sup>12</sup> Wonderly (2016, 232).

some philosophical object is close. Just as one might feel lost or alienated in the world without one's beloved, one might feel lost or alienated in the wake of losing a philosophical picture—in both cases, the bottom may drop out.

## §2.2 *Inaccessible value*

Consider now a very different kind of case, drawn from philosophical fiction. In a little-discussed example from a much-discussed paper, Frank Jackson imagines that we discover Fred, who “has better colour vision than anyone else on record; he makes every discrimination that anyone has ever made, and moreover he makes one that we cannot even begin to make.”<sup>13</sup> Jackson leverages this example in an argument against a metaphysical doctrine of physicalism, according to which “physical information” is the only sort of information that there is. But here I am interested in something else.

Suppose we agree with Jackson that “we should admit that Fred can see, really see, at least one more colour than we can.”<sup>14</sup> In addition to telling us purely descriptive information—for example, “There is an additional color” or “This object is not actually green”—Fred might make distinctly *evaluative* claims. Fred might declare, for instance, “You’re missing out on something extremely valuable.” The evaluative features of the new color, including how it compares to old colors, is as inaccessible to us as the new color itself. Nevertheless, if we share Fred’s aesthetic sense generally, then we may appropriately defer to Fred’s evaluative claim about the new color, and even hope or wish that we could see it for ourselves. In other words, we may both judge that we are missing something of value and regret that fact, largely on the basis of Fred’s testimony.

---

<sup>13</sup> Jackson (1982, 128).

<sup>14</sup> Jackson (1982, 129).

Inaccessible value only partly overlaps the case of grief. In both cases, those of us outside the relevant experience can rightly judge that there is something valuable to which we lack access (respectively, the additional color, or the person as known *via* a particular relationship). Moreover, under normal conditions it is appropriate for us to defer regarding many of the descriptive and evaluative features of what we do not experience or see. It seems reasonable to defer to the advanced perceiver about what the additional color is like, and it seems reasonable to defer to the appropriate griever about what the person was like. But in the case of color, interpersonal relationships play no psychological or normative role. It is just Fred's superior faculty of color perception that explains why he sees and appreciates what we cannot. There would be nothing inappropriate about us wishing to acquire an equally good faculty. On the other hand, it seems odd that we would wish to acquire all of the interpersonal relationships that make particular instances of grief appropriate.

Although one need not be in some interpersonal relationship to judge that one is missing something in the case of inaccessible value, acquaintance with the relevant phenomenon, in both cases, is plausibly a necessary condition for fully appreciating its absence. Only someone who first sees the new color, or knew the person, will really be able to imagine and feel the lack or loss—i.e., to *miss* having visual experiences of that particular character, or knowing the person.

In the case of the inaccessible value of a new color, something is objectively absent in the perceptual experiences of those who do not see it, but they are unable to directly and non-deferentially appreciate that absence. In the case of grief, something is merely subjectively lacking in those who are not in an appropriate position to mourn, due to the absence of a particular relationship rather than a perceptual deficit of some kind. These latter persons are, under normal conditions, still able to indirectly appreciate the loss by analogy to their own past or potential

grief. Deficient faculties are quite different from the mere absence of a particular interpersonal relationship, even though both may block direct access to some evaluative phenomena.

Philosophical angst shares with the case of color vision, but not grief, an objective purport; there is a positively important feature of the world that some people lack access to, but there is no barrier, in terms of their cognitive faculties, to their accessing it. But what the case of color vision itself lacks in terms of significance to life, both the objects of grief and philosophical angst have in abundance.

### §2.3 *Lost worlds*

Consider now certain problems involving what I will call “lost worlds.” What I mean by a “lost world” in this context is an all-encompassing (or, at least, much-encompassing) picture of the world and our place in it that is now largely inaccessible to us *qua* participants (rather than, say, *qua* sociologists or historians). The phenomena of nostalgia for a lost world and the fear of losing a world are closer to philosophical angst than grief or deficient color perception. To circumscribe the phenomenon, consider some remarks by Bernard Williams in *Ethics and the Limits of Philosophy*, regarding what he calls “real” and “notional” confrontations.

We should distinguish between *real* and *notional* confrontations. A real confrontation between two divergent outlooks occurs at a given time if there is a group of people for whom each of the outlooks is a real option. A notional confrontation, by contrast, occurs when some people know about two divergent outlooks, but at least one of those outlooks does not present a real option. The idea of a “real option” is largely, but not entirely, a social notion. An outlook is a real option for a group either if it already is their outlook or if they could go over to it; and they could go over to it if they could live inside it in their actual historical circumstances and retain their hold on reality, not engage in extensive self-deception, and so on. The extent to which they can do this depends on what features of their present social situation are assumed to remain constant if they go over to the other outlook (160-161).

What interests me here is how we would understand someone with whom we are in merely notional confrontation when they make distinctively evaluative judgments in favor of their own way of life. Let us just take Williams’ own example of a medieval samurai.

Many outlooks that human beings have had are not real options for us now. The life of a Bronze Age chief or a medieval samurai are not real options for us: there is no way of living them. This is not to deny that reflection on those value systems might inspire some thoughts relevant to modern life, but there is no way of taking on those outlooks. Even utopian projects among a small band of enthusiasts could not reproduce *that* life (161).

Suppose we recover a manuscript from the time period during which the way of the medieval samurai was on the cusp of extinction. In this imagined document, a samurai intensely laments the loss of his way of life on the grounds that it is immensely valuable and has no adequate replacement in the modern world. In other words, we discover the writings of an angsty samurai. If Williams is right that we have very little in the way of participatory access to the way of life that the medieval samurai represents, then it is plausible that we are incapable of fully evaluating the way of the samurai in a well-informed manner, and consequently of comparing it to our own way of life. In fact, supposing that what precedes the extinction of the medieval samurai is something to which we have no direct access, there are grounds for thinking that the samurai on the edge of extinction (when the new world is coming into existence) is *better*-positioned to make a comparative value judgment than we are; after all, he has had a taste of both worlds.

Even if we are moved by the thought that the samurai is in a superior epistemic position to us, I suspect that it is unlikely that we are actually prepared to defer to him about which way of life is better. At most we may endorse our way of life with a little more epistemic humility than we did before. Of course, seemingly *contra* Williams, many people take themselves to accurately and legitimately lament the modern world in favor of times past. But if these lamentations are fully intelligible, then it must be because these are cases where we *do* have some ability to appreciate what the respective past era was like. Perhaps it is possible to do this when there is sufficient sub-cultural overlap, for example, between the medieval scholastic and contemporary Catholic culture.

Lost worlds have much in common with both the cases of grief and color perception. But what is lost or absent is neither a matter of a special relationship nor a particular cognitive faculty; rather, what is lost or absent can be fully appreciated—for better or worse—as valuable only from within the relevant point of view, or form of life. Those outside are unable to appreciate the loss. And unlike the cases of grief and absent color, it is exceedingly unlikely that those outside will simply defer to those within. (Indeed, from the point of view of someone outside the angsty perspective, the loss might even seem like an improvement.)

The case of lost worlds also bears a similarity to grief that it does not bear to color. In the case of both lost worlds and grief, those of us on the outside may rightly judge that, from our point of view, there is (in a peculiar sense that may haunt us) no loss at all. A kind of relativism holds for both cases when it comes to the appropriateness of grieving or lamenting a lost world. As noted above, no such relativism exists in the case of color vision—failure to see the new color is simply a loss or regrettable absence for anyone who does not see it, whether they know it or not; there is no analogue for color perception being absent *for* one person who lacks it, rather than another.<sup>15</sup>

#### §2.4 *God revisited*

Finally, return to the paradigmatic philosophical case of angst, but which I noted above provides an exception to the contrast with grief: God. God's non-existence may cause angst for any number of reasons.<sup>16</sup> Perhaps the most popular reasons, at least in the philosophical

---

<sup>15</sup> Like grief over lost persons, grief over lost worlds might involve something like Wonderly's "security-based attachment" (see section 2.1). As with philosophical systems, I am not sure whether the world of the samurai counts as a "particular" in Wonderly's sense, but the other similarities between the issues involved are unmistakable.

<sup>16</sup> Several of these reasons (as well as reasons to prefer that God does not exist) are collected in Klaas Kraay (2018).

literature, have to do with an alleged role that God plays in the reality or objectivity of moral value.<sup>17</sup>

My purpose in this subsection is not to adjudicate the question of God's relationship to morality, but simply to illustrate an overlooked aspect of angst about the non-existence of God, which connects it directly to the case of grief. To those (especially Christians, Jews, and Muslims) who think of God as a particular person, imagining a world without God is much closer to imagining a world without one's beloved than it is to considering a more typical philosophical object. This similarity has little to do with morality, at least not in the direct sense owing to divine command theories of ethics. Consider that morality might be bunk, yet we might *still* lament the absence of those whom we love—including God. So, perhaps surprisingly, angst about God is separable from questions of morality, even though there are various ways of connecting the two domains.

Despite this similarity to grief, angst about theism nevertheless seems to be a distinctly *philosophical* kind of angst. Why? Is it just an accident of historically contingent intellectual taxonomy that beliefs about God remain under the purview of academic philosophy? I think not. For one, God does not “die” in the way that persons do—see the comments above on the difference between philosophical angst and grief. More importantly, in the case of grief, there is no pressure for *everyone* to lament a person's absence. Yet in the case of God, those who are angsty (who I will call “pro-theists”, following recent convention) characteristically join Nietzsche's madman to make a claim on every person's situation, effectively saying to every person, “You ought to grieve.” It is almost as if the death of *everyone's* beloved is being

---

<sup>17</sup> Probably the most classic statement of this appears in *The Brothers Karamazov*, in which it is claimed more than once that the non-existence of God entails that everything is lawful.

considered. It is both the case that what would be lost (or is already absent) is objectively good, *and* that lamentation is the objectively appropriate, universal response.

As an example of what this sort of angst-mongering looks like, consider how Blaise Pascal—famous for thinking of God in personal rather than philosophical terms—rails against those who purport not to care about whether or not God exists. For Pascal, God’s existence is the only bulwark against human misery.

This carelessness in a matter where they [people apathetic toward the existence of God] themselves, their eternity, their all, are at stake, irritates me more than it moves me. It astonishes and appalls me; to me it is monstrous. ... We do not need a greatly elevated soul to understand that there is no real and solid satisfaction here; that all our pleasures are only vanity; that our ills are infinite; and that finally death, which threatens us at every moment, must in a few years infallibly place us under the horrible necessity of being either annihilated or wretched eternally.<sup>18</sup>

There is a puzzle facing both Pascal and God’s lesser angst-mongers. In the case of theism, as with lost worlds and inaccessible value, those who are not angsty may suffer from radical epistemic inaccessibility with respect to the theistic possibility. While those who have either lived in the lost world (*viz.*, the world formed by the thought that there is a particular divine person) or somehow have other ways of appreciating it may find themselves unable to find any “real and solid satisfaction here”, the rest may find it easy to simply move on. Or, they may just continue living: there may be no sense of “moving on” at all. Though the angsty pro-theist might be unable to see her non-angsty compatriots as anything other than making a “monstrous” error in judgment, or missing something important about the world, there does not seem to be any sensible, rational way of inducing angst in those who lack it. After all, for someone who neither thinks that God exists *nor* appreciates what the theistic picture is really like, what motivation could there possibly be to get oneself to appropriately lament the non-existence of God? Why be a glutton for *that* kind of punishment?

---

<sup>18</sup> Pascal (2005, S681, 217).

Notice that, even though the lament over the imagined “death of God” shares with grief a sense that a particular person has been lost or is otherwise absent, only those who can appreciate the world in a particular way (as containing “the beloved”) can in turn fully appreciate the quality that the world had or can have with that particular person. In this way, pro-theistic angst is not unlike all cases of grief, but it is still quite unlike paradigmatic cases. The analogue in grief for Nietzsche’s madman is not simply someone who is trying to convince, say, a parent to mourn the death of her child; it is, rather, someone who is trying to convince a parent to mourn the “death” of a child who, by her lights, never existed in the first place. In another way, it is like trying to convince someone to mourn the death of a person they never knew, or a fictional character who, in some strange sense, might really have existed.

### §3     **Insiders and outsiders**

The above examples help to illustrate the way in which angst that is distinctly philosophical must have a universal purport. It offers a lament that purports to be both responsive to objective evaluative features of the world (say, that it *regrettably* lacks God) *and* appropriate for anyone who is attuned to the evaluative and other relevant facts. But those who want to communicate or induce angst that is distinctly philosophical face difficult epistemological and communicative problems, highlighted in some of the contrasts above. Those philosophical claims or viewpoints likely to induce angst are all- (or at least much-) encompassing. They contribute to how we see the world overall, as well as how we decide to live. These sorts of perspectives, and the value that they secure, can be difficult to appreciate from the outside. Even if those who know what it is like within the walls of, say, theism seem to themselves to *just see* that the world outside their windows is bleaker than the world within, this perception may be as inaccessible to outsiders as the alleged betterness of the samurai life is to all of us.

The example of the death of God, especially as it relates to grief, color perception, and lost worlds, suggests a further taxonomy of *insiders* and *outsiders* with respect to access to the value of some would-be philosophical truths. Insiders, whether they believe the thesis or not, have some kind of access to what it would be like both for the thesis to be true and for it to be false. An insider, in this sense, need not actually endorse the relevant thesis, although in many cases it might be a necessary condition that they at one time did endorse it. This necessary condition holds if endorsement is itself necessary for full appreciation of the relevant phenomena, but it will not hold otherwise. Put in the suggestive terminology of mourning, insiders are those who are acquainted, at least in thought, to the one whose loss or absence is lamented. Outsiders have access to at most what it would be like either for the thesis to be true or for it to be false, but not both. For some putative object of philosophical angst, knowing what it is or would be like, combined with the ability to picture its non-existence, is sufficient for being an insider.

Must outsiders be such that they *never* endorsed the relevant thesis? Put another way, does simply having been acquainted with the object of angst at one time or another, in reality or in mere thought, count toward being an insider? Although this is a tempting idea, it should be rejected. First, there may be merely nominal ways of endorsing a thesis that, due to their superficiality, fail to provide one access to a genuine sense of what the truth of the thesis is like, even though one still counts as endorsing the thesis. For example, it seems that there could be a nominal theist who is so unreflective and unattached to theism that she fails to gain any angsty attachment to God or the idea of God. Indeed, Nietzsche's madman can be interpreted as saying that his interlocutors in the marketplace are examples of this very phenomenon. Second, there is no obvious reason why someone who once endorsed a thesis might not experience philosophical forgetfulness, losing a sense of what it was like to stand within a particular point of view. Such a

person might have undergone a “transformative experience” in the sense developed by L.A. Paul, whereby not only her beliefs and perceptions have changed, but also her affective and evaluative apparatus.<sup>19</sup>

The epistemological upshot of the foregoing discussion is that the ability to convey appropriate philosophical angst to the angst-free might depend on various a-rational, contingent features of the non-angsty person’s history and psychology, namely the extent to which they have had access to, and recall, what some philosophical picture was really like from the inside. They must be insiders. This suggests that the only routes for convincing outsiders are, likewise, distinctly a-rational, for example, a form of radical conversion to a different, new perspective, or at least to a new vantage point, from which one can appreciate the evaluative contours of some philosophical contrast.

Recall a question from the previous section about the point of angst about God. I asked: what motivation could there possibly be to get oneself to appropriately lament the non-existence of God? The question generalizes: Even if one could make another person or oneself angsty on *any* particular issue, why do so? In fact, it is even open to question what the value is in expressing philosophical angst at all, whether or not one is able to convert others. In answer to the question about the point of making others or oneself angsty, the question is really only answerable from the perspective of seeing the angst as warranted or possibly warranted. From the angsty or angst-curious point of view, angst is or might be the fitting response to the world as it really is, including its dreadful or potentially dreadful aspects. It is plausible that there is intrinsic value in appreciating the world as it really is, rather than living under illusions, and so this intrinsic value is available as a reason to spread or acquire philosophical angst where it seems appropriate.

---

<sup>19</sup> See L.A. Paul (2016).

A second valuable feature of identifying where philosophical angst is appropriate has to do with philosophical inquiry itself. This also speaks to the second question about the value of philosophical angst in the event that no one else is convinced to acquire it. Philosophical angst is valuable, partly because it can direct and motivate our philosophical inquiry. It amounts to an appreciation of what is at stake in certain philosophical debates. That some questions are more angst-ridden than others is not a reason to pursue *only* the angsty topics, but it is a reason to *make sure to pursue* those topics, or to pursue them more centrally, and to never lose sight of them.

Finally, there is an important lacuna in what I have written so far. It is open to outsiders to think that it is *they* who are more in touch with evaluative reality than angst-mongering insiders. After all, it is possible that the world-picture that grips the insider is ultimately a confusion of some kind, perhaps requiring liberation *via* Wittgensteinian philosophical therapy. According to such a view, proper attunement to the evaluative and other features of the world results not in angst but in *freedom from angst*.

In response to this point, it seems to me that there is at least some very weak reason to suspect that angsty insiders are in a better position *vis-à-vis* delusions—namely, they have some taste of both worlds in question. This provides some grounds for a very mild sort of deference by outsiders toward insiders. However, the deference warranted is so weak that it may amount to no more than grounds for a brief second-guessing of one's own lack of angst. Moreover, even *if* angsty insiders are in a somewhat better epistemic position, it is still epistemically admirable that they second-guess their own angst, considering the possibility that they are in the grip of a confusion.

Ultimately, in the event that an insider and an outsider make parallel but competing claims about which of them is confused, it seems that a stalemate, or something very close to a stalemate, threatens. But this is not as worrying as it might seem to be at first, because it still seems valuable that the angsty and the angst-free continue to engage one another with an attitude of openness to the possibility that they are either missing something about the world or seeing something that is not really there.

#### **§4 Philosophical angst and contemporary philosophy**

A better understanding of philosophical angst can help to illuminate underexplored aspects of contemporary debates, as well as to make intelligible the underlying concerns that animate those debates. In this final section I give two such examples: the metaphysics of reasons and the meaning of life.

##### *§4.1 The Metaphysics of Reasons*

Within metaethics, perhaps the angsty philosopher *par excellence* is Derek Parfit. Parfit defends an externalist (in his taxonomy, “objectivist”) account of reasons, according to which the reasons that we have are both irreducibly normative and non-natural. What makes Parfit angsty, however, is not just the normative metaphysic that he happens to endorse, but the importance that he attributes to the truth of that metaphysic. According to Parfit, if his account of reasons is false, then nothing matters.

Unsurprisingly, several philosophers—including philosophers otherwise sympathetic to Parfit’s metaethical program—have balked at the dramatic inference Parfit draws from the falsity of his view. The point here is not to consider these responses in detail, but to notice a structure that they share.

Julia Driver, Larry Temkin, and Sharon Street all agree that Parfit is onto something when he says that his opponents' views would mean that nothing matters. Yet they each agree, for related reasons, that this fact is not as significant as Parfit thinks it is. Driver confesses that she shares with Parfit "the same worries that anti-realist views of morality ... are at least *prima facie* disturbing," though she does not quite say why.<sup>20</sup> Temkin writes that "Parfit may be correct that if his externalist non-natural view of normativity is mistaken, then nothing matters in his sense. And I, for one, *would very much like it to be true*, and *would be very disappointed* if Parfit's view was false, so that nothing mattered in his sense."<sup>21</sup> Like Driver, Temkin does not explain *why* it would be so disappointing to him if Parfit's view were false, other than to say that it would mean nothing matters *in Parfit's sense*. Finally, Street concedes to Parfit that, on anti-realist views like hers, "Nothing matters, ultimately, independently of the attitudes of beings who take things to matter."<sup>22</sup> For her part, Street registers no disappointment. Despite some concessions, each one of these philosophers argues that Parfit ignores *intermediary ways* that things can matter, which are established by non-nihilistic alternatives to Parfit's externalism.

Both Driver and Street endorse versions of metaethical constructivism that deliver some of the features of reasons that make Parfit's view so attractive. For example, Driver observes that one appealing feature of Parfit's non-naturalism is that it secures the non-contingency of moral norms. But some forms of constructivism might also secure a kind of non-contingent universality of reasons—at least, *for beings like us* (namely, in Driver's theory, beings who are "moral agents and evaluators"). Though Driver concedes that this sort of view does not "solve"

---

<sup>20</sup> Driver (2017, 172).

<sup>21</sup> Temkin (2017, 27, emphasis mine).

<sup>22</sup> Street (2017, 121).

the contingency problem (presumably, because reasons are still contingent at the level of agent-type), it “mitigate[es] it.”<sup>23</sup> In other words, one feature of reasons that plausibly contributes to Parfit’s angst seems to come in degrees. Once we recognize this, we can see how the metaphysics of reasons is not a zero-sum game. Some anti-Parfitian views accommodate these features better than others.

In the same spirit, although he does not name a specific view, Temkin outlines “intervening positions” between Parfit’s non-naturalism and normative nihilism. Each of these positions accommodates varying degrees of “robustness” as a property of reasons—where robustness is a gloss on non-contingency.<sup>24</sup> So, although Temkin prefers that Parfit’s non-naturalism be true, he disagrees with Parfit that we face a situation of non-naturalism or bust.

With an account of philosophical angst in hand, we can interpret Parfitian angst in a way that puts some pressure on the objection from intermediate positions. According to a version of Parfit’s view understood in terms of philosophical angst, the non-naturalistic, purely and irreducibly normative nature of reasons is *existentially important* and *irreplaceable* by other possible natures that reasons might have. Instead of asking (with Parfit) whether Parfit’s interlocutors have the concept of a reason, or asking (with Parfit’s opponents) whether there are intermediate kinds of mattering that Parfit ignores, we can ask what it is about external reasons and externalist mattering, on Parfit’s conception, that is so important and irreplaceable. Parfit’s opponents see his angst as a response to some property that comes in degrees, whereas it is clear that Parfit himself thinks there is something at stake in an absolute, all-or-nothing sense.

---

<sup>23</sup> Driver (2017, 184-85).

<sup>24</sup> Temkin (2017, 29-30). Temkin models robustness with differing font sizes in writing the word “REASONS”, illustrating that there are many font sizes between the large font representing Parfitian reasons and the very small font representing nihilism.

This is a fruitful way to understand a position like Parfit's, because it forces us to ask more specifically what the particular contribution is of this or that philosophical scheme to what matters to us in the relevant domain. In another paper, I argue that what is angst-inducing about non-naturalistic moral realism is that only on such a view does moral justification (and so, moral vindication) enjoy a kind of robust stance-independence that is qualitatively, rather than quantitatively, different from the many degrees of independence secured by different versions of moral anti-realism.<sup>25</sup> A similar route may be available to the Parfitian: perhaps *only* non-naturalistic, purely and irreducibly normative reasons are sufficiently independent of our psychologies to get us what we (or, at least, what *insiders*) want (that is, what we want at an existential or emotional rather than a theoretical level) out of an account of reasons.

If this interpretation of Parfitian angst is correct, then all Driver and Temkin have shown is that there is a kind of non-contingency that comes in degrees of distance from actual, individual attitudes, but which is still tethered to them. They see Parfit's position at the most non-contingent end of this spectrum, and nihilism just over the edge of the least non-contingent end. Hence, the more plausible views we can develop in the middle, the more of a buffer we have against nihilism. It is an intuitive thought, to be sure. But, for Parfit, *the entire kind* itself--the kind of non-contingency that comes in degrees, that is--is qualitatively different from the kind of non-contingency secured by his own view of reasons.

#### §4.2 *The Many Meanings of Life*

The topic most closely associated with philosophical angst is probably the meaning of life, and in particular the relationship that the meaning of life bears to the existence of God. There is a common thought that the existence of God is a necessary condition for the meaning of life.

---

<sup>25</sup> See chapter 3.

Justifying the common thought is a lot trickier than merely expressing it. But, again, the point here is to show how the concept of philosophical angst might illuminate some underlying concerns of this debate, not how it might solve it.

Due especially to the work of Susan Wolf, it is common to distinguish between at least two distinct senses of “meaning” that are sufficiently similar to provide a plausible and important disambiguation in the common thought.<sup>26</sup> On the one hand, we can think of there being a kind of cosmic or ultimate meaning, which only God, or something like God, can secure. This is the sort of meaning that attracts intensification or capitalization, as in *real meaning* or Meaning. The concept of ultimate meaning is somewhat hard to articulate any further. What, after all, could such meaning amount to? Candidates include things such as *being the result of divine action*, or *being created by someone for a purpose or plan*, or *having an objective purpose*, or perhaps simply *being an object of divine love*. However we define the sort of meaning that only something godlike can provide, philosophers generally agree that it is distinguishable from a kind of meaning that life might have irrespective of whether God exists. This is the sort of meaning often exhibited by, for example, valuable personal projects like painting a picture, raising a child, climbing a mountain, or proving a complicated theorem in mathematics. Even in a world in which lives lack ultimate meaning, these activities remain more *meaningful* than, say, counting ants for no reason.<sup>27</sup> So, it seems that there is an index of something appropriately called “meaning”, even

---

<sup>26</sup> See Wolf (2010).

<sup>27</sup> Robert Paul Wolff conveys the following anecdote about the eminent myrmecologist E.O. Wilson: “After the usual greetings, he showed me the centerpiece of the office, a large table on which, under a Plexiglas dome, was a bustling, complex ant colony. Wilson banged the side of the table, which set the ants scurrying, and as they poured out of the anthill he pointed out the soldier ants, worker ants, and so forth. I didn't have much in the way of conversation. What can you say about an anthill, after all? So, casting about for something to say, I mused aloud, ‘I wonder how many ants there are in the entire colony.’ ‘Fifteen thousand,’ Wilson replied. ‘How can you be sure?’ I asked. ‘I counted them,’ he said.” (Wolff, “What Have I Been Reading,” September 1, 2013,

though these activities do not seem to provide whatever we are talking about when we talk about ultimate meaning.

This distinction between cosmic and ordinary meaning (what Wolf calls meaning *of* life and meaning *in* life, respectively) is sufficient for deflating (to some degree, at least) the common worry about the existence of God and meaning. Even if God is required for cosmic meaning, we may still have (ordinary) meaningful lives even if God does not exist--and that is enough. What else might the pro-theist say?

The pro-theist may argue further that ultimate meaning is *existentially important*, rather than merely good or desirable; it is not, like winning the lottery, something that would be very good but unnecessary for a good life. Second, the pro-theist may argue that ultimate meaning is *irreplaceable*, in that the need that it satisfies is not also satisfiable by something else, even ordinary meaning.

This is a fruitful way to understand the angsty concern about God and meaning, because it invites us to ask more specifically what the particular contribution is of this or that kind of meaning to our lives going well. In another paper (and in section 2.4, above), I argue that what is properly angst-inducing about, at least, popular forms of personalistic theism is that God is a particular person to whom one might stand in a meaningful relationship, akin to other personal relationships but with deeper and more lasting consequences.<sup>28</sup> Hence, unlike in the case of metaethics, I do *not* think that the standard form of philosophical angst here is ultimately cogent: there seems to me to be no sense of cosmic meaning the angsty recognition of which is

---

<http://robertpaulwolff.blogspot.com/2013/09/what-have-i-been-reading.html>). My brother, who is a myrmecologist only slightly less eminent than Wilson, claims to have counted 20,434 ants while performing research at the Konza Prairie Biological Station at Kansas State University. Both myrmecologists presumably counted their ants for reasons, though it is open to question whether these reasons are sufficient for meaningfulness.

<sup>28</sup> For elaboration, see my "The Value and Particularity of Persons" (ms).

incumbent upon *everyone* who considers the non-existence of God, given that God is or would be a particular person; like grief, angst about God requires a particular relationship. But that is not the point here. I am arguing that those philosophers who *do* want to defend this kind of angst may fruitfully adopt the framework here on offer. Their explanatory burden is to specify, first, why cosmic meaning is existentially important in human life; and second, why it is irreplaceable by ordinary meaning with respect to this role.

## **§5 Conclusion**

This paper is as an exercise in the phenomenology of a characteristic attitude, one that I call “philosophical angst.” But it also provides a framework for exploring new avenues in extant philosophical debates. In service of the first goal, I have identified the conditions and normative status of philosophical angst, comparing and contrasting it with overlapping but distinct attitudes that are not properly philosophical. In service of the second goal, I have suggested two applications of the notion of philosophical angst, one in the debate about the nature of reasons, and one in the debate about the meaning of life.

## Chapter 2: Metaethical Value Judgments

Suppose that moral nihilism is false. Does anything of further, evaluative importance depend on which non-nihilistic metaethical theory is true? This paper defends metaethical claims (including non-nihilistic metaethical claims) as intelligible and appropriate objects of evaluative judgments. I will call such a judgment a “metaethical value judgment”: an evaluative judgment about ways that the world might be, metaethically speaking. Such evaluative judgments are distinct from descriptive judgments about what some theory can or cannot accommodate. It is one thing to judge that a metaethical theory, say, cannot account for moral error, but quite another to judge that it would be worse if the theory were true (though it very well might be worse *because* it cannot account for moral error).<sup>29</sup> Thinking about metaethical value judgments affords us a useful, clarifying way of thinking about what is at stake in metaethics. If some metaethical value judgments are cogent, then metaethics is not an evaluatively inert domain, even if the traditional conceptual distinction between meta- and normative ethics might suggest as much.

This chapter gives some general, abstract argument for thinking that there are significant evaluative differences between metaethical possibilities, and it responds to some reasons for skepticism toward the general project of making metaethical value judgments. This chapter sets

---

<sup>29</sup> Sturgeon (1986) addresses the question of what “difference” it makes whether moral realism is true, but his examples involve *desiderata* that are non-evaluative in my sense, e.g., accommodating the possibility of moral error.

the stage for making a substantive case about what we should desire and hope for when it comes to metaethics.

## §1     **Metaethics matters: some negative arguments**

The distinction between meta- and normative ethics gives us some reason to doubt the claim that there is something evaluatively at stake between non-nihilistic metaethical theories. With the exception of moral nihilism, all metaethical theories provide *vindications* of the moral domain. That is to say, each non-nihilistic metaethic, were it true, would establish how and why moral discourse and practice are legitimate and well-grounded. In this sense, a true, non-nihilistic metaethic will in turn vindicate whatever the best normative moral theories happen to say, even though the true metaethic will not itself dictate what they say. So, if it is normative (including applied) ethics that makes a difference to what is evaluatively important in morality, it seems that there cannot be much at stake evaluatively between non-nihilistic metaethical theories. Sure enough, morality itself is at stake in the question of nihilism. But despite their many differences, non-nihilistic metaethical theories all have the same ultimate evaluative upshot: the vindication of normative ethics and whatever the best normative theory happens to say.

Before addressing this argument head-on, it is worth acknowledging that there is, in fact, not a clear or non-contentious dividing line between meta- and normative moral theory. Here is a simple gloss on the intuitive distinction: metaethics concerns the nature and status of the moral domain—that is, its metaphysics, epistemology, semantics, and psychology. Normative ethics, on the other hand, concerns the content of the moral domain—that is, its constitutive norms. But there are views that seem to violate this conceptual boundary. Some versions of metaethical constructivism, according to which moral truths are determined or constituted by some actual or

hypothetical agreements among moral agents of suitably specified capacities, are hard to place in the traditional taxonomy, since they seem to do double-duty as meta- and normative ethical theories. Likewise, some metaethical theories seem especially well-fitted to particular normative theories—for example, the naturalistic versions of moral realism developed by Richard Boyd and Peter Railton seem to comport especially well with some form of normative consequentialism. Nevertheless, I think that the simple gloss tracks a real conceptual distinction, even if some extant views in moral philosophy do not obey that distinction and must have their individual meta- and normative components carefully parsed.

Fortunately, there is a response to the intuitive argument against the evaluative significance of metaethics that does not depend on calling into question its distinction from normative ethics. From the fact that metaethics is not normative ethics, or that evaluative facts are determined only by the latter, it does not follow that metaethics is of no evaluative significance. Notice that theories about God or freedom of the will are not instances of normative ethics, yet no one doubts that competing theories in those domains have significantly different evaluative upshot. Of course, unlike theories of God or freedom, metaethical theories are explicitly vindicatory of the normative moral domain, and so the very domain of evaluative significance itself seems to depend on one such theory being true. But it *still* does not follow that metaethical theories have no evaluatively important differences. First, two metaethical theories might share their *moral* evaluative upshot, but differ in their *non-moral* evaluative upshot, which is not dependent on (moral) metaethics. For example, even though two metaethical theories equally vindicate the claims of morality, perhaps only one of them is congenial to the demands of prudence, or with concerns about meaning, or with aesthetic norms.

It can be easy to miss, but many seemingly non-evaluative objections to metaethical theories can be interpreted or developed in just this direction. For example, some philosophers object to metaethical non-naturalism and Platonism on the grounds that such theories make the epistemology of morality too mysterious. There are two ways to interpret this objection. As stated, it is often a purely descriptive sort of objection. It points to a theoretical desideratum, *viz.* accommodating moral knowledge, and says that some theories fare poorly. But there is a nearby evaluative version of this complaint: *it would be bad* if we lacked all access to moral truths, properties, and the like. One need not think this condition would *morally* bad, and certainly not *only* morally bad, in order to endorse this evaluative complaint. Such a result would, at the very least, be *epistemically* bad.

Second, and most importantly, we should not be deceived by the intuitive theoretical hierarchy that moves from applied, to normative, and finally to metaethics, even if we are talking about the broadest version of such a hierarchy: the study of normativity in general. The fact that metaethics takes normative ethics as its object, theoretically speaking, does not mean that normative ethics cannot in turn take metaethics (and everything else) as its object, normatively speaking. Normative evaluation has effectively unlimited scope: everything can come under our normative evaluation.<sup>30</sup> In a word, suppose that some metaethical theory vindicates the normative domain. This vindication will include what the normative domain says about the metaethical theory itself.

---

<sup>30</sup> In my view, this is one of the lessons of G.E. Moore's Open Question Argument. According to that argument, of any non-moral property or thing putatively identified with goodness, we can coherently and competently ask, "Is it good?" Even if this argument does not accomplish everything that Moore intended, the persistence of "open questions" is well-explained by the thesis that normative evaluation has unlimited scope. Cf. Robert Adams' claims in *Finite and Infinite Goods* that there is a "critical stance" intrinsic to morality that allows any *natural* or *empirical* fact or phenomenon to fall under the critical purview of morality (78).

The above considerations are negative arguments against the intuitive idea that metaethics does not matter evaluatively, provided that some form of non-nihilism is true. Although I regard these arguments as sufficient for motivating an exploration of what is evaluatively at stake in metaethics, the motivation would be bolstered by a positive argument for thinking that metaethics *does* matter in the relevant sense.

## §2     **Metaethics matters: a positive argument**

Provided that we countenance metaethics as a genuine sub-domain of moral inquiry, we can consider the following argument, which I will call the “Evaluative Difference Argument”, in favor of the claim that there probably are evaluative differences between different metaethical theories.

1. The moral domain is valuable and important.
2. The major families of metaethical theory offer radically different descriptions of the moral domain.
3. The fact that two descriptions of something valuable and important are radically different provides some defeasible evidence that they are not equivalent in evaluative upshot.  
Therefore,
4. The fact that the major families of metaethical theory are radically different provides some defeasible evidence that they are not equivalent in evaluative upshot.

In a sentence, the intuitive thought behind this argument is this: when there are two very different descriptions of something that is extremely valuable and important, it is somewhat likely that it makes a significant difference which description is true. Before defending the details of the argument, I will offer two clarifications about the premises. By “major families of metaethical theory” I mean those theories that are typically distinguished at the highest levels of abstraction, warranting, for example, different categories in philosophical encyclopedias and different units in introductory course syllabi. So, for example, Richard Boyd and Peter Railton, though their metaethical views are not identical, offer theories in the same family: naturalistic

moral realism. Likewise, despite their differences, Alan Gibbard and Simon Blackburn defend two theories in the same family: metaethical expressivism. The difference in family is what makes the difference in theories “radical”: theories from different families are generally incompatible at a fundamental, or near-fundamental level. Unlike many of their individual components, they are mutually exclusive. Both expressivism and naturalistic realism cannot be true, even though some components of expressivism (e.g., non-cognitivism about moral judgment) are compatible with some components of naturalistic realism (e.g., objectivism about moral truth).

By “equivalent in evaluative upshot” I mean equivalent with respect to how good or bad it would be if the theory were true, and in what ways. Setting metaethics to one side for a moment, two incompatible theories of God are evaluatively equivalent, in my sense, if it would not matter, evaluatively speaking, which theory is correct. That it does not matter evaluatively which theory of God is correct does not mean that the debate in general does not matter: after all, it may be an intrinsically interesting debate, or it may be important for some contingent, instrumental reason. But the absence of evaluative difference does mean that the debate is, in itself, evaluatively inert. On the other hand, if two theories of God are *not* evaluatively equivalent, then there must be *some* respect in which it is correct to say that the truth of one theory would, in itself, be *better* than the other. Just as we might naturally say “It doesn’t matter which of X or Y is true” regarding evaluatively equivalent theories, a natural way of expressing evaluative non-equivalence is, “It would be better if X rather than Y were true.”

So much for the meaning of the premises. I will now give some reasons for accepting premises 2 and 3. For present purposes, I am just taking premise 1—that morality is valuable and important—as uncontroversially true. The argument is simply not aimed at those who do not care about morality at all, or regard it as having no importance or value. Rather, it is aimed at

those who care about morality, but do not think that much hangs on metaethics—provided some non-nihilistic theory is true. Notice that premise 1 is weaker than a nearby claim that only some philosophers would accept but many would deny, *viz.*, that the moral domain is *maximally* valuable and important, or enjoys *overriding* value and importance. The weaker version should be acceptable even to philosophers who find themselves in the tradition that dethrones morality from its maximally authoritative perch.<sup>31</sup> The view that moral demands do not necessarily settle practical matters should not be conflated with the view that moral demands are not important at all, or even that they should not factor heavily in practical deliberation. As far as I know, even philosophers who devote considerable time to dethroning morality do not reject premise 1.

Premise 2 is more controversial than premise 1, but I do not think that it should be *very* controversial. Different families of metaethical theories ground moral discourse and practice, not to mention moral truth itself, in totally different phenomena. It is difficult to overstate the difference between, say, non-natural, irreducibly normative facts, on the one hand (in which robust realists ground morality), and our deepest commitments and cares, on the other (in which expressivists ground morality). That being said, it must be admitted that differences between the families of metaethical theories come in degrees: non-naturalist realism is further from naturalist constructivism than is naturalist realism. This point about degrees will be relevant in defending and better understanding premise 3.

Premise 3—the fact that two descriptions of something valuable and important are radically different provides some defeasible evidence that they are not equivalent in evaluative upshot—is the most controversial premise in the Evaluative Difference Argument. The first

---

<sup>31</sup> Some important works in this contemporary tradition include Bernard Williams (1985), Harry Frankfurt (2004), and Susan Wolf (1982).

thing to emphasize is that the premise makes an evidential or probabilistic claim, not a claim of entailment. So, counterexamples to the stronger entailment claim would not by themselves defeat the probabilistic claim. How might the evidential claim be defended? By thinking about examples that are less taxonomically complicated than metaethics (which itself overlaps several areas of philosophy), I think that Premise 3 can be, at the very least, made more intuitively plausible. Consider, again, two radically different theories of God: the traditional, supernaturalist view that God is a particular person who created, governs, and loves the world, and the more “philosophical” view that God is the “ground of Being,” beyond and wholly separate from worldly categories like “person”, “creator”, and certainly “love.” It would be extremely surprising if the debate between these two views turned out to be evaluatively idle—that is, if there were no cogent reasons whatsoever to think that it would be *preferable* that one rather than the other theory were true. This surprise is well-motivated by the fact that radically different descriptions of something extremely significant are likely to yield different evaluative upshot.

If the simple example of concepts of God is compelling, then the case of metaethics should be even more compelling. Whereas theories of God are narrowly focused on the nature of an important particular being or concept, metaethical theories are sprawling accounts of an entire important domain of human inquiry, practice, and (one hopes) knowledge, themselves spanning metaphysics, epistemology, semantics, and psychology. The chances that two radically different theories in such a domain would be evaluatively equivalent is exceedingly low.

That being said, recall what I said above about degrees of difference between theories. Differences between the families of metaethical theories come in degrees, our expectation that something evaluative hinges on which of two theories are true should somewhat track how similar or different those theories are from one another. The Evaluative Difference Argument is

extremely general, but it should not be interpreted as saying that, once radically different theories are on the table, something of equal importance is at stake between all of them.

Finally, it is important to appreciate the defeasibility of the kind of evidence defended by the Evaluative Difference Argument. There will be cases where two theories are radically different, but not in a respect that matters evaluatively. For example, numbers are valuable and important. Not only are they intrinsically interesting, but they allow human beings to understand and exercise considerable control over the world. Consider nominalism and Platonism: these are radically different accounts of numbers. According to Platonism but not nominalism, mathematics is grounded in actual Platonic entities, which have a kind of robust being. Yet, it is not clear that there is something *evaluatively* at stake between nominalism and Platonism. The Evaluative Difference Argument is compatible with this observation. As far as the argument says, we had *prima facie* reason to wonder whether something was evaluatively at stake between nominalism and Platonism. If we discover that there is nothing evaluatively at stake here (say, because everything important and valuable about mathematics is preserved by both theories), then we have simply defeated the evidence that the Evaluative Difference Argument points to.

### §3     **Some puzzles about metaethical value judgments**

The Evaluative Difference Argument provides some reason to think that metaethical disputes, even between rival non-nihilistic theories of morality, are not themselves evaluatively idle. So, the project of identifying and defending the correct metaethical value judgments is a live one. But there are some reasons for thinking that the project is confused. First, unlike theories about God or freedom, metaethical theories themselves interpret value judgments, judgments like *it is good if moral realism is true* or *it would be better if moral facts were socially constructed*. This gives rise to a worry about there being a merely verbal dispute between philosophers who

offer competing metaethical value judgments from incompatible metaethical points of view. Second, given that metaethical theses purport to be necessary, there seems to be something odd about saying that it would be better or worse if one rather than another theory were true, since at least one of them could not be true at all. This second problem generalizes to any evaluative judgment about a putative metaphysically necessary truth, and so what I say here is of general interest beyond metaethics. Finally, in the special case of nihilism, there seems to be something even more problematic about saying that it would be bad or worse if nihilism were true, since if nihilism were true, *nothing* would be bad or worse—there would not even be a moral domain to theorize about. In the remainder of this section, I develop and address these three problems.

### §3.1 *The first puzzle: contested terminology*

It might seem that metaethicists cannot intelligibly hold fixed the meanings of words like ‘good’ and ‘bad’ when evaluatively contrasting metaethical possibilities, raising the specter of a merely verbal dispute. In contrast, in judging that it would be bad if atheism were true, one is not employing any terminology that is contested from within the debate about theism. Theists and atheists can hold fixed the meanings and referents of words like ‘good’ and ‘bad’ in order to proceed in a substantive evaluative disagreement about the evaluative upshot of the theistic possibility. But metaethical theories themselves give interpretations of the evaluative terms of the debate.

For example, suppose that a moral realist says of the quasi-realist possibility, “It would be bad if quasi-realism were true.” The quasi-realist may deny this, asserting instead, “It would not be bad if quasi-realism were true.” While the two metaethicists make apparently contradictory statements, do they have contradictory meanings? Is the realist really asserting something that the quasi-realist is denying? You might think not, if the right way to read the assertions involves

interpreting the moral predicates as they are interpreted by the speakers' theories.<sup>32</sup> If that is the right way to read such assertions, then the above exchange can be disambiguated as follows:

Realist: It would be bad<sub>REAL</sub> if quasi-realism were true.

Quasi-realist: It would be not be bad<sub>QUASI</sub> if quasi-realism were true.

This looks a lot like the disambiguation of a merely verbal dispute, and that would be bad news for the project of making and defending metaethical value judgments. While such judgments would not be wholly unsalvageable, they would merely become fodder for debates between people with shared metaethical commitments.

Notice first that, *if* it is right to read the dispute as merely verbal, at least for the reasons given, then it does not seem that speakers who endorse incompatible metaethical theories could ever successfully engage in any evaluative debate. On such a view, only speakers who endorse equivalent (or consistent) metaethical views could intelligibly disagree about what is valuable. But this is simply untenable. It is manifest that a quasi-realist and a moral realist can disagree about, say, whether the death penalty is morally permissible, without equivocating between permissibility<sub>REAL</sub> and permissibility<sub>QUASI</sub>. And what goes for any ordinary first-order value judgment should go for a first-order metaethical value judgment, even if the latter is somewhat exotic. There is no reason to think that the specter of a verbal dispute arises *only* when the subject matter is metaethics rather than something else.

The fact that the meaning of metaethical value judgments are not affected by the speakers' respective theories does not mean that what is actually true metaethically makes no difference to nature of the judgments, however. In thinking about how metaethical truth affects the meaning of metaethical value judgments, there are two main routes for interpreting the disagreement: we

---

<sup>32</sup> See Francén Olinder (2012, 2016) for an account of "metaethical pluralism," according to which the truth conditions of a speaker's moral assertions partly depend on her metaethical standpoint.

can read the evaluative predicates as moral or non-moral. To illustrate the first strategy, let us (for reasons that will become clear momentarily) restrict our attention to non-nihilistic possibilities and consider an example metaethical value judgment, with the reading of the evaluative terms made explicitly moral:

[1] It is very morally bad that moral realism is false and quasi-realism is true.

Notice that there is no obvious reason why [1] cannot be true, even though its propositional form implies that quasi-realist theory is in fact correct. Suppose, for example, that quasi-realism really is true. The correct way to understand the meaning of [1], in that case, is as an expression of a suite of deep desires, commitments, longings, and the like, that take the moral domain, realistically construed, as their object. No incoherence arises.

Even anti-realist theories that do not accommodate talk of truth and moral properties, like a simplistic version of emotivism, may still give a semantically vindictory reading of the sentence. On such a view, asserting [1] may simply be a way of expressing disapproval or dislike toward the fact that anti-realism rather than realism is true, as in, “Realism—hoorah! Anti-realism—boo!” And disagreements over [1], if such a view is true, amount to exchanges in which speakers try to push each other into having their respective attitudes toward metaethical possibilities.

Now consider a different assertion, which does not, in the way of [1], imply that some anti-realist theory is in fact correct.

[2] It is very morally bad if anti-realism rather than moral realism is true.

Suppose that [2] is asserted by someone in a state of uncertainty about metaethics, for whom both realism and anti-realism are epistemically possible. Then, from the point of view of the speaker, the best reading of the moral predicate in [2] is simply left open. Here there is no need to

worry about the implications of assuming that realism is actually false. Again, no incoherence arises.

In looking more closely at the dialectical possibilities, the sense of a puzzle begins to fade away. What became of the appearance of a verbal dispute? The best way to understand evaluative debates between speakers who endorse incompatible metaethics is in terms of a shared evaluative space on the first-order level, *about which* there is a theoretical disagreement. In this way, evaluative disputes about metaethical possibilities are no different from other evaluative disputes. Provided that none of the speakers believe that the truth of their metaethical theory, but not their interlocutor's, would *entail* the truth of a disputed value judgment (in which case problems of circularity threaten), the differing metaethical accounts of the evaluative terminology does not pose a special problem. Metaethical value judgments are again shown to be unproblematic despite being somewhat exotic.

### §3.2 *The second puzzle: metaphysical necessity*

Suppose we compare two metaethical possibilities and agree that at most one of them is metaphysically necessary, and at least one of them is impossible. After all, metaethical theories are partly metaphysical in nature, and metaphysical theses characteristically purport to be true of metaphysical necessity. This might seem to cause trouble for comparative value judgments about the two theories, since we are guaranteed to be evaluating an impossibility. Likewise, it would be odd to ask whether it would be better if 10 were prime; it is not even so much as logically possible that 10 is prime, and so you might think that it also would not be anything, evaluatively speaking.

First, it should be controversial whether metaethical truths are metaphysically necessary. There are reasonable doubts about necessity in normative theory, and these might apply to

metaethical theory as well.<sup>33</sup> It is conceivable that a metaethical theory might be contingently true. As an example, consider the following reading of J.L. Mackie's argument for nihilism in *Ethics: Inventing Right and Wrong* from the strangeness of moral properties. According to Mackie, perhaps there could be such moral properties, for example, "to-be-pursuedness," but they are so strange that they cannot be countenanced in *this* (allegedly) naturalistic world. Yet the world *might have been* non-naturalistic, and as such have accommodated moral properties. Indeed, Mackie seems to have thought that the existence of God could make some kind of difference to whether the universe can countenance moral properties. But (so the argument goes) since the world is not that way, we should not believe that there are moral properties.<sup>34</sup>

I suspect that some will be convinced that metaethical theories might be contingent, but many will not be. So, suppose for the sake of argument that metaethical theories are the sorts of things that, if true, really are metaphysically necessary. Even if metaethical value judgments are unintelligible in reference to *known* metaphysical impossibilities, they still might be intelligible in reference to *epistemic possibilities*. After all, this is a reasonable way to understand familiar evaluative worries about other theses that purport to be necessary, for example, theism.<sup>35</sup> As long as the existence and non-existence of God are epistemically possible for us, it is sensible for us to wonder about which possibility would be better or worse.

---

<sup>33</sup> For some of these doubts, see Gideon Rosen (ms).

<sup>34</sup> I have not found many who endorse this reading of Mackie, though one prominent example is Richard Joyce (2016a; 2016b, 192 fn. 11). Cf. Jamie Dreier, "Mackie's Realism: Queer Pigs and the Web of Belief."

<sup>35</sup> A widely shared dogma in philosophy of religion says that, if God exists, then God exists necessarily. See Swinburne (2004, 79, 148) for doubts and Leftow (2010) for discussion.

Perhaps evaluative worries about a known metaphysically necessary or impossible theory really are confused or incoherent.<sup>36</sup> But it should be controversial whether metaethical value judgments are unintelligible even in reference to known metaphysical impossibilities. Some philosophers argue—plausibly, in my view—that we can intelligibly distinguish between different impossible worlds, as opposed to the view either that there are no impossible worlds, or that there is just one impossible world in which all propositions are true.<sup>37</sup> If so, then we should be able to distinguish between these worlds evaluatively as well. And if we can do that, then we can plausibly distinguish between different, even impossible, metaethical worlds evaluatively. For example, just as it seems that, *if the number 10 were prime, then it would not be divisible by two*, so it seem that *it would be bad if the number 10 suffered constantly and for no reason*. If this sort of approach is right, then the fact that some metaethical theory is impossible may turn out not to prevent the us from intelligibly making a value judgment about its being true or false.

So the necessary purport of metaethical theses turns out not to be a problem for making metaethical value judgments. There is not a cogent worry about verbal dispute, and at the very least, such judgments make sense so long as their objects are epistemically possible. Moreover, if recent work on distinguishing impossible worlds is right, such judgments can even make sense when they regard known metaphysical impossibilities.

### §3.3 *The third puzzle: evaluating nihilism*

Recall that in response to the first puzzle, I bracketed the question of moral nihilism. I did this because reading the evaluative predicates in metaethical value judgments as moral poses

---

<sup>36</sup> Cf. Kahane (2016): “If nihilism is false, then we could in principle have reasons to fear it, if we don’t know it’s false. But such reasons are odd since we could have them only if the object of our fear isn’t realized. And it’s anyway hard to see why we would have these reasons if nihilism couldn’t make things worse for us or anyone.” Kahane does not distinguish between contingent and necessary version of nihilism, as I do above.

<sup>37</sup> See Daniel Nolan (1997) and Linda Zagzebski (1990) for defenses of impossible worlds.

special problems if one of the possibilities we are considering is the nihilistic possibility. The puzzle is that, if nihilism were true, then nothing would be good or bad. Moral nihilism, unlike quasi- and anti-realism, entails the falsity of all first-order moral judgments-- including metaethical value judgments.

The best response to this puzzle is to interpret the evaluative predicates with contents not contested by moral nihilism. The judgment that *it is bad if moral nihilism is true* must not receive a moral interpretation for its evaluative predicate. If the value in question is not moral, then the third puzzle evaporates. There is no problem if the value that nihilism denies is not the kind of value predicated of metaethical possibilities.

Of course, this solution--and much else that I have said so far--only works if we do not believe, and are not considering, *global evaluative nihilism*, or what is sometimes called *normative nihilism*. If we believe that global evaluative nihilism is true, then I doubt that there remains anything about which we can intelligibly ask evaluative questions. The only appeal in that case might be to impossible worlds, although even that would suffer from semantic mysteries about how we would get a purchase on evaluative terminology at all, in a globally nihilistic context. Consider: If some causal theory of reference is true, for example, the popular theory due to Kripke and Putnam, then it might be that persons within even merely contingently nihilistic worlds cannot refer to moral properties, though such properties are possible. They would be like Putnam's brains in vats, whose lack of referential access to brains and vats precludes them from even so much as formulating the hypothesis that they are brains in vats. In our context, this view would yield the parallel consequence that there can only be contingent nihilists in worlds where contingent nihilism is false. This is because contingent nihilism is impossible to believe or even contemplate without referring to moral properties, and, according

to the causal theory of reference, such reference is impossible in a world without such properties. (If this is right, then the existence of contingent nihilists would be decisive evidence that contingent nihilism is false!)

To sum up, the only way that the world might be that would undermine the making of metaethical value judgments is globally nihilistic. Even local forms of nihilism are compatible with the making of metaethical value judgments, since the evaluative predicates in such judgments may receive interpretations outside the domain in which the local form of nihilism is true.

#### **§4 Conclusion**

In the first two sections, I showed that the value and importance of morality combined with the radically different accounts of it given by competing metaethical theories provides us with some evidence that these theories have different evaluative upshot. And in the third section, I dispelled doubts about the coherence or sense of making metaethical value judgments in the first place. If these two arguments are successful, then we have a satisfying way of insisting that metaethics matters. It is not the case that what matters evaluatively is one thing, and metaethics another. Rather, metaethical possibilities themselves fall under the purview of normative philosophy—we may evaluate such possibilities just as we do possibilities in religion, philosophy of mind, or anything else, and we should expect that the possibilities are quite different in evaluative upshot.

Value questions about metaethics, then, belong alongside other examples of what Guy Kahane calls “the value question in metaphysics.”<sup>38</sup> And since morality is a very, if not the most,

---

<sup>38</sup> Kahane (2012).

important domain of human discourse and practice, it belongs alongside the other so-called “big questions” not just conceptually, but in importance.

## Chapter 3: Moral Realism and Philosophical Angst

As discussed in Chapter 1, philosophers express angst about the existence of God, the immortality of the soul<sup>39</sup>, freedom of the will, the nature of mind<sup>40</sup>, the metaphysics of modality<sup>41</sup>, and other subjects. On such issues, some of us do not only try to figure out what is true, we *worry* about it. Consider the case of theism. Positions range anywhere from the view that life is meaningless without God, to the view that it is terrible with God.<sup>42</sup> Less attention has been devoted to similar questions in metaethics. Yet, just as someone might not only wonder, but worry about whether theism is true—so they might not just wonder, but worry about whether a particular metaethical theory is true. Given the immense importance of morality, it is somewhat surprising that questions of angst in this area of inquiry have been neglected. This paper, I hope, gives *angsty metaethics* the attention that it deserves.<sup>43</sup>

---

<sup>39</sup> Eli Hirsch (2009, 2, 20) says that death and radical deception are “ineffably horrible” and “undermine everything that I have ever cared about in my life.”

<sup>40</sup> Jerry Fodor (1989, 77) says that “if it isn’t literally true that my wanting is causally responsible for my reaching, and my itching is causally responsible for my scratching, and my believing is causally responsible for my saying.... If none of that is literally true, then practically everything I believe about anything is false and it’s the end of the world.”

<sup>41</sup> Robert Adams (1979) argues that modal realism justifies “moral indifference.”

<sup>42</sup> See the essays recently published in Kraay (2018).

<sup>43</sup> See Guy Kahane (2012)’s pioneering work on what he calls the “value question in metaphysics.” Although Kahane is more narrowly interested in evaluative judgments rather than full-fledged philosophical angst, this paper could be read partly as a metaethical contribution to the general research program advanced by Kahane.

After sketching a general account of philosophical angst, I describe and defend a particular instance of it: angst regarding moral realism. Following the convention in the literature on the value of God's existence, in which "pro-theist" is used to refer to someone who thinks that it is good that God exists (whether or not they believe that God actually exists), I will use "pro-realist" to refer to someone who hopes that moral realism is true and worries that it might not be (whether or not they believe that realism is true). According to pro-realists, the truth of realism is extremely important, and there are no adequate theoretical alternatives: it is realism or bust. The phenomenon of pro-realism is immediately recognizable, but it is less obvious how one might understand and defend it. After all, as long as moral nihilism is false, how could something of great normative significance hang on which non-nihilistic metaethic is true? This paper seeks to both make angst about moral realism intelligible and to commend one way that pro-realists can make their case to the angst-free. That being said, I do not pretend to have exhausted the considerations that might motivate (or defeat) philosophical angst about moral realism.

## **§1 Philosophical Angst**

Perhaps we know philosophical angst when we see it, but it is useful to have a more exact account in order to isolate angst-inducing phenomena in any particular domain. Here is how I understand the phenomenon: Philosophical angst is a complex cognitive judgment combined with a conative component. The judgment is that something existentially important is also irreplaceable. *Existential importance* is positive importance to the fundamental meaning, intelligibility, or overall value of our lives and the world. If you think that a claim might be false but do not think that its truth is important in this sense, then you do not count as angsty. *Irreplaceability* means that there is no available surrogate for whatever plays the important role in question. If you think that an existentially important claim might be false but you also think that

some alternative claim would do just as well, then you do not count as angsty. The conative component, licensed by the judgment, is a kind of anxiety that what is important and irreplaceable in this sense is unavailable. Notice that, to be philosophically angsty, one need not believe the claim that one is angsty about, but both the claim and its falsity must at least be live options.<sup>44</sup>

It is no surprise that the most typical catalysts for philosophical angst are subjects including God, freedom, and death, rather than the ontology of numbers or the semantics of counterfactuals. The former but not the latter subjects are especially significant in how we think about the meaning of our lives. But this explanation in terms of existential significance makes it at least somewhat surprising that *metaethics* is not subject to more angsty reflection than it is. After all, metaethics concerns the metaphysics, epistemology, semantics, and psychology of arguably the most important domain of value that there is--and minimally, a very important one.<sup>45</sup> As a candidate for philosophical angst, metaethics seems intuitively closer to the philosophy of religion than to the ontology of numbers.

The metaethicist who most readily comes to mind as angsty is probably Derek Parfit. In *On What Matters*, Parfit famously (to some, notoriously) declares that, if metaethical non-naturalism is false, then nothing matters.<sup>46</sup> This declaration satisfies the cognitive conditions for philosophical angst, provided that it includes or presupposes the thought that it is existentially

---

<sup>44</sup> For a full account of philosophical angst, see Chapter 1.

<sup>45</sup> See my “Metaethical Value Judgments” (ms) for a general argument that we should expect different metaethical theories to have significantly different evaluative upshot.

<sup>46</sup> “If there were no such facts [facts that only “irreducibly normative claims” could state], and we didn’t need to make such claims, Sidgwick, Ross, I, and others would have wasted much of our lives. We have asked what matters, which acts are right or wrong, and what we have reasons to want, and to do. If Naturalism were true, there would be no point in trying to answer such questions. Our consolation would be only that it wouldn’t matter that we had wasted much of our lives, since we would have learnt that nothing matters” (Parfit 2011, 367).

important that something matters, and that no rival theory secures an adequate replacement for non-naturalist mattering. But Parfit's angst is narrower than it appears at first glance. He thinks that views rival to his collapse into nihilism or something very close to nihilism. And Parfit does not offer a substantive proposal for why things matter only on his metaethical view; rather, their mattering only on his view is an immediate logical or conceptual consequence of the ways that he has (somewhat idiosyncratically) defined the relevant terminology and concepts, as several commentators have pointed out.<sup>47</sup> Because he thinks that his rivals' views collapse into nihilism, Parfit's angst really amounts to the worry that nothing matters if nihilism is true—a claim much less controversial than it seemed at first. Moreover, it is not even clear that Parfit thinks it would be a bad thing if nihilism were true. Indeed, in volume 3 of *On What Matters*, he distances himself from the claim that it matters whether something matters.<sup>48</sup>

These limitations of Parfit's angst provide a good foil for illustrating some of the central concerns of this chapter. In contrast to Parfit, we might be angsty about metaethics not because we think that most views collapse into nihilism. We might instead think that, despite the fact that many views do not collapse into nihilism, they nevertheless fail to accommodate something irreplaceably important to the meaning or intelligibility of our lives. This latter kind of angst permits more philosophical charity, because it evaluates theories on their own terms, allowing that they basically accomplish what they purport to. In this chapter, I am wondering about metaethical angst in this richer, more charitable sense. While I will discuss realism's relationship to moral nihilism in its own right, I will not only discuss that; I am also interested in how realism compares to its non-nihilistic rivals—even when they succeed by their own lights. Compare: in

---

<sup>47</sup> See Street (2017), Temkin (2017), and Driver (2017).

<sup>48</sup> See n. 14.

thinking about freedom of the will, someone might be angsty if, in their view, only libertarianism about freedom secures something (in particular, libertarian freedom) that is both existentially important and irreplaceable. But this is substantively different from thinking that only libertarianism secures freedom in any sense. Someone (even a libertarian) may grant that compatibilist theories genuinely secure freedom of the will, but nevertheless still insist that such theories do not secure something extremely valuable or existentially important about the nature of freedom of the will, as it is described by libertarianism.

In short, angst about moral realism, or pro-realism, is the view that the truth of moral realism is necessary for something irreplaceable and existentially important to our lives going well, combined with the worry that realism might be false. This does not mean, however, that the pro-realist must think that all alternative theories collapse into moral nihilism.

## §2     **How to be angsty about moral realism**

In this paper, I am concerned with moral realism in a fairly robust sense. According to moral realism in this sense, there are objective moral facts, and these can serve as the content of true beliefs in the same way that characteristically descriptive non-moral facts may do so (for example, the facts that *human beings evolved* or that *two and two make four*). We are able to know and express moral reality in propositional form, and we have, by-and-large, made genuine progress in moral inquiry. The moral facts are both necessary and radically response- or stance-independent, in that they are not fundamentally grounded in human minds or attitudes, whether individual, collective, actual, ideal, or otherwise hypothetical.<sup>49</sup> Pro-realism says that something about this picture irreplaceably plays an existentially important role in the meaning or

---

<sup>49</sup> Notice that I have said nothing about the naturalism/non-naturalism distinction, even though “robust” moral realism is often associated with the latter. I am unbothered if my characterization of realism entails or otherwise requires non-naturalism, but I am also not committed to its doing so.

intelligibility of our lives, and so the fact that realism might be false warrants anxiety. It is not possible here to canvass all instances of non-realism in their dizzying variety, so in what follows I will highlight three general contrasts: the contrast between realism and moral nihilism, between realism and anti-realism, and, finally, between realism and quasi-realism.

The reasons for judging that it would be bad if moral nihilism were true are distinct from the reasons for judging that it would be bad if any form of non-realism (including both anti- and quasi-realism) were true. Unlike moral nihilism, anti- and quasi-realist theories are what I will call *vindictory metaethics*, in the sense that they (by their own lights) justify moral discourse and practice. In short, they are not self-conceived as error theories. The pro-realist, however, is not only concerned that morality be given a positive account, *à la* non-nihilism (a hope that anti- and quasi-realists may also satisfy), but that it be given a distinctly realistic account. It is not just morality *per se* that the pro-realist thinks enjoys existentially important irreplaceability, but morality realistically construed.

### §2.1 *The badness of moral nihilism*

Consider a simple example of an unmitigated injustice. In the *History of the Peloponnesian War*, Thucydides gives us an exchange between some Athenian invaders and their victims, the Melians. The Melians have no chance of resisting conquest. In emphasizing the pointlessness of resistance, the Athenians accurately observe that “the strong do what they can and the weak suffer what they must.”<sup>50</sup> What does this have to do with metaethics? If moral nihilism is true—that is, if morality is bunk—then there is no moral problem with the aforementioned relationship between the strong and the weak. The dominance of the strong over the weak may be painful or regrettable in some non-moral sense, but it would not be morally

---

<sup>50</sup> Thucydides (1996, 5.89, 416).

unjustified or otherwise morally objectionable. This is a good ground for angst about the prospect of moral nihilism, because it seems *bad* that there would be no grounds for objecting morally to this pitiful state of affairs. Not *morally* bad, of course—after all, nothing morally substantive follows from the truth of moral nihilism. But bad in some kind of non-moral sense; our lives (or, at least, the lives of the weak) are worse off if they lack grounds for moral objection. Accordingly, it seems appropriate to hope that morality is not bunk.

This simple argument is intuitive, but is the hope really reasonable? In what non-moral sense is it “bad” to lack moral justification? As mentioned, it cannot be that moral nihilism results in some morally bad situation, since nihilism results in no morally valenced situations at all. Moreover, morality does not itself prevent the strong from doing what they can or the weak from suffering what they must; morality does not change the world. If morality does not change the world, then it is hard to see how nihilism would be bad even in a prudential or instrumental sense. What good does it do us that we can legitimately morally object to something, if it is just going to happen anyway?

Nowadays it is common to distinguish between multiple domains of value, including but not limited to morality, meaning, aesthetics, and self-interest. One’s life might be good along some dimensions but not others. A predominantly self-interested person might live a life abounding in prudential value (say, they always get what they want) but sorely lacking in moral value. Likewise, the “moral saint” might live a life abounding in moral value but sorely lacking in meaning.<sup>51</sup> The acknowledgement that moral value does not exhaust the values that we should (and do) care about is an advancement beyond a myopic focus on moral value. However, it is also important to see that moral value may *contribute* to, or even be required for, the fulfillment

---

<sup>51</sup> Susan Wolf, “Moral Saints.”

of other kinds of value. It is true that a person's life may be abounding in moral worth yet lacking in meaning, but it is also true that a life entirely lacking in moral worth may—in virtue of that lack—suffer from meaninglessness. In answering the question about what kind of badness accrues to us if nihilism is true, the pro-realist must appeal to one of the non-moral domains of value to which moral value nevertheless contributes. The domain of meaning is well-suited to this purpose.

What morality does, *vis-à-vis* the Athenian-Melian dynamic, is secure for us a particular normative standing in virtue of which we enjoy an authoritative claim against injustices. It is much better that we have this kind of standing and claim, even when we have no way of enforcing it. It is good—even non-morally good—to be morally worth treating well. Unjustified harms and oppression cause one's life to go poorly, but one's life is better if one has available an authoritative protest—even when it is ineffective. Such a protest secures meaning in the midst of harms, in that it makes one's condition intelligible as morally unjustified, and objectionable on these grounds. There is no adequate, non-moral replacement for this standing and claim if morality is bunk. At best, the death of morality leaves in its wake various instrumental and epistemic norms, but these are not even inconsistent with the aims and values that the Athenians press against the Melians. Moral standing differs from other kinds of standing not in degree, but in quality. Over and above instrumental, epistemic, and other norms, moral standing provides the strong with a powerful reason not to harm the weak.

This defense of angst about nihilism locates something that morality delivers for us and appeals to its irreplaceability and importance in our lives. It is worth noting that this view is consonant with some provocative remarks by Thomas Nagel regarding the importance of inviolability in any account of moral rights. Nagel writes, "What actually happens to us is not the

only thing we care about: What may be done to us is also important, quite apart from whether or not it is done to us—and the same is true of what we may do as opposed to what we actually do” (108). Again, “not only is it an evil for a person to be harmed in certain ways, but for it to be permissible to harm the person in those ways is an additional and independent evil.” And finally,

We can distinguish the desirability of not being tortured from the desirability of its being impermissible to torture us; we can distinguish the desirability of not being murdered from the desirability of our murder's being impermissible; we can distinguish the desirability of not being coerced from the desirability of its being impermissible to coerce us. These are distinct subjects, and they have distinct values. To be tortured would be terrible; but to be tortured and also to be someone whom it was not wrong to torture would be even worse (111).

Being the sort of being whom it is permissible to torture is bad, even despite the fact that it is in another way worse to be “killed unjustly than, say, accidentally.” Nagel’s point is not that permissible torture or murder are themselves worse than their impermissible counterparts; rather, it is “being someone it is not wrong to torture” that is worse than being someone it is wrong to torture—whether or not one is tortured.

Objection: What if the realist world is such that, even if it has moral value, it is, on the whole, very bad for us—for example, if it is full of little more than intense and constant suffering, and the best normative theory says that the suffering is morally unjustified? Does the non-moral goodness of mere (violated) moral status really outweigh or otherwise trump such a state of affairs? Surely, a realistic possibility is only better than its nihilistic counterpart if the net value in the world, for us, turns out positive. In contrast to a very morally bad realistic world, an otherwise good or even so-so nihilistic world seems better for us. If moral nihilism is true, then what is otherwise unjust suffering is not morally bad at all; it is nothing, morally speaking. And it is plausible that a set of morally bad states is worse than a set of states that are neither morally

good nor bad. So, perhaps it is not true as a general proposition that it is better if realism rather than moral nihilism is true.<sup>52</sup>

In response to this objection, consider an implication of the view it commends. If an otherwise morally bad state is better if it is neither morally good nor morally bad, then finding out that one's otherwise unjust suffering is morally neutral should be encouraging news. If a would-be morally bad state (say, the state of being betrayed) was better on moral nihilism, then finding out that moral nihilism is true should be a kind of appropriate comfort for a person experiencing betrayal. Such a discovery would not alleviate a person's anger or unhappiness, but it would mean that they were not morally wronged, and that no particular reactions are morally justified.

But this appeal to moral nihilism for comfort is wrongheaded. It is better not only that there is moral value; it is better even that morally bad things really are so. If one is deciding between having a set of morally neutral things (say, trips to the beach) or a separate set of moral evils (say, trips to the torture chamber), then, all else equal, it is reasonable to choose the neutral things. If the question of preferring realism over nihilism were analogous to such a decision, then it would be reasonable to prefer nihilism in some of these cases. It is from this fact that the objection derives its plausibility. But this is not the right way to think about comparing metaethical possibilities. The would-be goods and evils must be held fixed in the comparison: *this*

---

<sup>52</sup> This sort of thought motivates Kahane (2016). As noted above, it is also given expression at various points by Parfit in Vol. 3 of *On What Matters*, e.g., "It is a difficult question whether and how it matters whether anything matters. If we believe that suffering matters greatly, we may regret this fact. We might try to believe that, as Nihilists claim, nothing matters, because we have no reason to care about anything. We might then conclude in despair that Nihilism is false, because some things, such as suffering, really do matter." Again, "I don't know whether I would be ... very disappointed if I came to believe that nothing matters in [the] reason-involving sense. I am not glad, for example, that suffering matters. But since I believe that we have reasons to care about suffering, and that we have other, weaker reasons to care about some other things, I am trying to understand these reasons better." Unlike Parfit, I am glad that suffering matters.

*bit of would-be mercy, this instance of would-be wrongful treatment*, etc. We are considering not just whether we have certain goods or evils period, but what it would be like if *the very things* we already rightly believe to be good or bad (like torture) were not really so.

Contrasting moral nihilism and non-nihilism bears some similarity to contrasting first-order normative theories that have different evaluative upshots, with the aim of figuring out which would be better. When contrasting Kantianism and Utilitarianism, for example, it is not dialectically appropriate to radically change what the world is like when one imagines the truth of the rival theory. In wondering whether it would be better if Kantianism rather than Utilitarianism is true (as Nagel does), it is not appropriate to simultaneously toggle whether or not pleasure and pain states exist. Rather, one holds fixed the world as it otherwise is—including the existence of pleasures, pains, and rational capacities—and imagines the implications of Kantianism and Utilitarianism for *this* world.

To illustrate the general view on offer with an example, not only is torturing innocent uncles for fun morally bad, but we rightly do not want it to be otherwise. The question of whether torturing for fun is good or bad is not itself an evaluatively idle question. It is good for uncles, and for us, that it is impermissible to torture them for fun. This higher order desire is itself normatively justified, in addition to the justification for the first-order view itself. A world in which torturing innocent uncles is morally permissible, let alone good—holding everything else fixed—is worse than one in which it is morally evil. That is to say, it is to some extent better that torturing innocent uncles is morally bad rather than morally good.

The intuition that the permissibility of torture would itself be bad is related to a popular normative objection to simplistic versions of divine command theory, according to which one ought to do whatever God happens to command. The objection says that such a theory allows

God to make it the case that actually evil actions are good, and that this counts against the appeal of the theory. Returning to the Melian dialogue, the point here is not that the Melians should be happy that their moral claim and status are violated, nor should they wish that their violation was in fact permissible. Rather, the point is that they should prefer having a moral claim and status, given the facts that actually constitute the violation, to not having them.

Ultimately, one need not be convinced by the details either of Nagel's view or my view regarding, respectively, what is important about either a deontological first-order theory of morality or a non-nihilistic metaethic. *Whatever* one thinks are the goods (including non-moral goods) conferred to us by morality, *these* are vindicated by any non-nihilistic metaethical theory. In other words, take whatever goods you think are delivered by the best first-order moral theory. These goods are, in turn, vindicated by the true non-nihilistic metaethic, and it is to these that we can refer in justifying an angsty aversion to the possibility of nihilism. In order to reject this argument, you must think that morality makes no net-positive non-moral contribution to our lives. If, *per impossible*, we could choose whether to live in a morally nihilistic or non-nihilistic world, and could not appeal to moral values in making the decision, the person who rejects the argument of this section must say that we have no positive reason to choose the non-nihilistic world. And that, I think, is very implausible.

## §2.2 *The badness of anti-realism*

To consider moral nihilism is to consider the possibility that there are none of the goods of morality; accordingly, these goods are dialectically available for licensing evaluative judgments about the possibility of moral nihilism. Moral anti-realism is trickier, because to consider anti-realism is to consider whether there *are* a great number of the goods of morality, but understood

in an anti-realistic way.<sup>53</sup> Speaking loosely, angst about moral nihilism is concern about the *existence* of morality, but angst about anti-realism is concern about the *nature* of morality.

The rest of this subsection applies to anti-realist theories other than quasi-realism, due to special complexities arising from the pan-expressivist or minimalist semantic program partly constitutive of quasi-realism. In going expressivist or minimalist at every level of discourse about the moral, the quasi-realist makes it especially difficult to say how quasi-realism differs from realism in its depiction of the moral domain, which in turn makes it difficult to say what is at stake between the two theories. Because this issue arises for none of the other varieties of anti-realism, it is useful to categorize and treat them separately.

So what *is* angst-inducing about (non-quasi) anti-realism? Why might it induce angst in someone if the moral badness of, say, murder is not realistically construed? Why might it induce angst in someone if moral facts are, say, constructions grounded in collections of subjective attitudes? Or if moral truths are contingent commitments that human beings have endorsed over time, for evolutionarily-specified purposes? Or if moral facts are, fundamentally, facts about the results of idealized decision-procedures among fully-informed agents?

Standard, longstanding disagreements between realists and anti-realists would suggest answers having to do with objectivity, non-contingency, and the like. I suspect that there is something potentially regrettable in the anti-realistic interpretation of each of these aspects of morality, but the answer that I want to pursue here has to do with what I consider to be the radically independent character of moral justification if moral realism is true. Following Sharon

---

<sup>53</sup> Note on terminology: There is an unfortunate terminological parallel between “pro-realism” and “anti-realism.” With “anti-realism” I refer to anti-realist metaethics, not an angsty aversion to moral realism. In Chapter 4, I employ the only slightly happier term, “anti-moral realism”, to refer to the angsty aversion to moral realism.

Street, I regard some kind of mind- or response-dependence of moral phenomena to be constitutive of moral anti-realism. Here is Street:

According to the anti-realist, if an agent has normative reason to X, then this conclusion must somehow follow from within her own practical point of view: if the conclusion that she has reason to X is not entailed from within the standpoint constituted by her own set of evaluative attitudes, then she does not have that reason.<sup>54</sup>

As Street elaborates in a footnote, “the point of contention between realists and anti-realists about normativity is the answer to the central question of Plato’s Euthyphro (in rough secular paraphrase): are things valuable ultimately because we value them (anti-realism), or do some things possess a value that holds independently of us and our attitudes (realism)?”<sup>55</sup> Street locates her own view on the anti-realist side of this “Euthyphronic” account of the debate. (Notice that Street’s own exclusion of quasi-realism in this part of the discussion is what permits the Euthyphronic account, since quasi-realists emphatically deny that things possess moral value in a response-dependent way.)

Street’s exact characterization is a bit too narrow for my purposes, since it seems to count idealizing theories as realistic. It is also too focused on an agent’s normative reasons, rather than moral value and truth generally. The version of the realist/anti-realist distinction that will I utilize here counts our hypothetical, idealized, and collective responses as part of “us and our attitudes.” So, here is the modified definition of anti-realism that I will employ in what follows:

According to anti-realism, moral truths hold in virtue of us and our attitudes, broadly construed: if a putative moral truth is not grounded in our evaluative attitudes—whether actual or idealized—then it is not a moral truth.

Following this taxonomy, if moral realism is true, then it is not just the case that our moral reasons are *external* in the traditional sense of being object-given, or intrinsic to the situations or

---

<sup>54</sup> Street (2016, 295).

<sup>55</sup> Street (2016, fn. 9).

objects that our reasons are about or for—after all, these very claims can be true *in virtue of* facts about moral agents *qua* valuing beings. And it is not just the case that moral properties are this way—after all, even a simple version of subjectivism can consistently claim that torture is wrong because it causes unjustified pain, not because anyone thinks it is wrong. It is the *metaphysical grounds* of such truths that are at issue. On realism, we may have moral reasons, and things have moral properties, in virtue of a system of moral truths that is itself independent of persons and their attitudes. Various metaethical parties offering positive accounts of morality may agree to some version of the claim that, for example, the suffering and oppression of the weak generates an object-given reason for us to care and do something, and that these reasons are reasons to stop something that is itself morally bad whether or not anyone thinks so. But there is an additional kind of independence involved in the true moral verdict: morality itself, the system of moral truths itself, declares in favor of the oppressed, ruling their oppression a normative impossibility—with no reference to human moral attitudes in the grounds of this impossibility. When we think rightly, we concur with this verdict that did not come from us, either literally in our needing to issue it, or metaphysically in our attitudes, independent or collective, being needed to ground it.

Is angst on this basis intelligible? One way of beginning to see that it is, is to appreciate how realistic independence is *qualitatively* rather than *quantitatively* different from anti-realistic independence. There is, the pro-realist can concede, a kind of anti-realistic independence that is a matter of degree or interpersonal distance, which I will call the quantitative sense of independence. Consider, for example, a view on which moral claims are true in virtue of their being the subject of hypothetical agreement among idealized agents. The pro-realist should not deny that there is a recognizable kind of independence involved in such an account. After all, this

kind of view explicitly distances moral truths from the actual responses of individuals. But notice that the view does not make moral truth at all independent of individual responses in what I call the qualitative sense. After all, on this kind of anti-realism, the responses of each person's idealized self are partly constitutive of the moral truth, and so the truth is still tethered to responses. The same goes for other response-dependent theories. For instance, any view on which moral claims are true in virtue of one's robust, consistent attitudes over time, is a view on which there is quantitative independence, but it is independence of the anti-realist kind that is explicable in terms of degrees of distance from the actual responses of individuals. Anti-realist moral truth is still tethered to actual, individual responses.

In being qualitatively different, realistic independence is not just one more step on the ladder up and away from the actual attitudes of individuals. From the pro-realist point of view, for the system of moral truths to be *really* independent is for it to be independent of our attitudes in *any* sense--wholly untethered from actual, individual responses. This is what I mean by the "radical independence" of moral truth. After all, even the rational agreements among our idealized selves are subject to the strictures and demands of a realistic morality. From the realist point of view, rational agreement among idealized selves might always produce true judgments, but such success is understood as the success of tracking, rather than constituting or constructing, the independent truth.

So far, I have made some effort to characterize the difference between realistic and anti-realistic independence, but I have not yet attempted to explain what would be so *good* (and irreplaceable) about realistic independence. First, a caveat: someone who finds moral realism of the sort described to be itself unintelligible will struggle to see how there can be anything other than degreed (and hence tethered) independence from individual responses. In a way, such a

person is not the target of my arguments, any more than the person who claims not to understand the concept of God is a target of arguments that it would be good if God exists. On the other hand, I do want to show that, given that someone thinks that there is such a possibility, it is at the very least rationally intelligible that they would not find anti-realist independence a fitting surrogate for realist independence. Fitting surrogates must be relevantly similar in kind to their originals, and I aim to show that anti-realist independence is not such a surrogate for realistic independence.

The pro-realist claims that the radically independent nature of moral justification is good, but not merely good; it plays an existentially significant role in the meaning and intelligibility of our lives.<sup>56</sup> This is because it militates against a kind of normative lonesomeness or moral bootstrapping, in which—in one way or another, at bottom—we are responsible (either causally or constitutively) for providing our own moral vindication.<sup>57</sup> Again, I allow that pro-realist angst might be only intelligible from within the perspective of someone who “gets” what it would be like to see morality in a realist way, in which context radically independent moral justification is

---

<sup>56</sup> Here it is worth noting that, in a continental context, my account of why anti-realism rightly induces philosophical angst is closest to the account of angst (or anxiety, or dread) one finds in Jean-Paul Sartre’s *Being and Nothingness* and “Existentialism is a Humanism.” Sartre locates the origin of anxiety in the fact that human beings are radically responsible not only for their own actions but for the very principles that guide them—in the sense that there is no objective, independent moral reality that determines this in advance. Of course, Sartre did not have in mind contemporary metaethical categories and one cannot know whether he would have also rejected anti-realism, but my account of what is lost without realism is recognizably closer to his than it is to, say, the systems of Søren Kierkegaard or Martin Heidegger, who have quite different concerns when they speak of angst.

<sup>57</sup> Cf. the following extended excerpt from Wright (1995: 226), who makes a related point (though one whose import he ultimately rejects) about external sanction: “In general ... the immediate price of anti-realism about morals is merely[!] that the gravity of moral judgement will lack an external sanction. When one is asked, ‘Why bother to try to arrive at correct moral opinion?’, the only available answer will be: because such an opinion informs *better* conduct - better, that is, from a moral point of view. The value of moral truth will thus be an instrumental, moral value. It is common to think that there are, by contrast, intrinsic, general values associated with pure discovery, understanding and knowledge of the real world. Properly to characterise and to understand such values seems to me to be a very difficult task. In any case, for the moral anti-realist, that kind of value cannot attach to moral truth. But I think it has seemed important that it should only because of the tendency of philosophers to suppose that there is nothing for truth to be that is not associated with value of that sort.”

available. But against the backdrop of such radical independence, there seems to be little difference, *vis-à-vis* bootstrapping, between the different degrees of quantitative, tethered independence available on anti-realism. Relative to a realist framework, increasing the degrees of tethered independence is like increasing the length or complexity of a bit of circular reasoning for a desired conclusion.

Here is an elaboration of the analogy to circular reasoning. Imagine that you take yourself to have an independent (now, in the sense of *non-circular*) argument for an existentially important claim (say, that human beings are worthy of respect). You then encounter a group of theorists (these are the analogues to anti-realists) who give a variety of circular arguments for the important claim. If circular reasoning is all these theorists think there could ever be, then, perhaps, length and complexity is the most you could want or hope for in terms of “independence”. But once one sees—or takes oneself to see—an independent, non-circular route to the important claim, it makes sense that the circular routes, including the longest and most complex among them, will pale in comparison. Anti-realistic justification is to realistic justification as circular justification (for an important claim) is to non-circular reasoning. Although anti-realistic justification is not formally circular, contrasted with realistic justification, it is as disturbingly *self-referential*.

More can be said to defend the intelligibility of pro-realist angst. Although the considerations that follow do not depend on the analogy to circular reasoning, they are similar in spirit. Perhaps ironically, the perspective on offer here is sympathetic to (but not quite the same as) an argument due to the quasi-realist Simon Blackburn, who is hardly a friend of the sort of robust moral realism in consideration here. Blackburn argues that several anti-realist views face a problem of moral relativism. For example, Blackburn offers the following objection to the

constructivist metaethic associated with Christine Korsgaard, according to which each person is rationally bound to self-legislate moral norms in virtue of an inescapable practical identity.

Korsgaard's view, he says, fails to preclude "pluralities of self-legislating persons whose identities are happily bound up in various constraints they set themselves under, but who unfortunately find these constraints in entirely different places."<sup>58</sup> How does this relate to pro-realism? The pro-realist may adapt Blackburn's worry into the present discussion by saying that what would be bad or regrettable about Korsgaardian constructivism—even if it turned out to be true—is precisely that, between any two different persons or communities of persons who locate normative constraints "in entirely different places," there is no radically independent source of authority to adjudicate the dispute, to vindicate either one side or none. This is reminiscent of the concern about moral bootstrapping that I expressed with an analogy to circular reasoning: if everyone can appeal to "constraints they set themselves under", then they are ultimately justified self-referentially, by concerns that others need not share.

To illustrate the point by another route, consider yet another challenge by Blackburn, this time aimed squarely at the moral realist. Blackburn demands that the moral realist explain what *more* it is important to "say" about moral evils beyond the moral condemnation to which anti-realists are equally entitled. "What is wrong with [the Taliban in its relation to women] is that the men oppress the women, impoverish their lives, and keep them in a state of ignorance and inactivity. Why should we feel any urge to say more than that? Isn't it bad enough?"<sup>59</sup> The pro-realist should reply: the urge is not that there be something else that we might say, but that there

---

<sup>58</sup> Blackburn (1999: 219). It must also be noted that Blackburn criticizes several theses associated with realism on the same grounds: in particular, McDowell's sense theory and neo-Aristotelian accounts that ground morality in a conception of human flourishing (218-220).

<sup>59</sup> Blackburn (1999: 223).

be more than just “we” saying it. The traditional realist answer was that the Taliban is not just wrong, but objectively and robustly so. The angsty realist does not deny this. But the really important point is that it is not the *saying* of such a thing that matters. It is rather that there is an independent verdict issued, as it were, by reality itself. What is “said” remains the same. Earlier I said that true moral statements “concur” with an independent verdict. The point is that true moral assertions are repetitions, restatements of a reality already partly constituted by the truths that the assertions express, much like *two and two make four* or *no sphere has corners* are reiterations of radically response-independent realities.

Finally, it is worth considering a somewhat peculiar objection. The objection is worth considering not because it is necessarily powerful in itself, but because it helps to illustrate one further (and final) way of bolstering the intelligibility of pro-realist angst. Some readers might suspect that there is a kind of crypto-theism lurking in the background of the perspective just advanced, and they might want to launch a debunking argument against pro-realism on that basis. Isn’t pro-realist just a watered down version of angst about the “death of God”, who provided a kind of robust, external reference point for morality and meaning? In a way, the suspicions is right, but it provides confirmation of the overall framework defended here. Think of the familiar existential desires associated with theism: to “have God on our side”, or not to be “alone in the universe,” or to achieve and to see “cosmic” justice.<sup>60</sup> These desires plausibly derive from fundamental religious attitudes for which realist value is also fitting. Incidentally, they also provide additional support for David Killoren’s recent argument that robust moral realism is a kind of “religion.”<sup>61</sup> Even if this connection may in turn provide grounds for a kind of debunking

---

<sup>60</sup> The terms in quotation marks are pulled from the ether, not any particular source.

<sup>61</sup> David Killoren (2016).

argument against a certain style of religiously-infused moral theorizing, that is a matter of what is *true*, not what would be *better*--and irreplaceable.<sup>62</sup> (It would, of course, be disappointing for pro- realists if the very grounds of their angst could be employed as defeaters for realism itself!) Notice further that the sorts of replies that the angst-free may give to pro-realists are manifestly unsatisfying in their analogue religious form. To the pro-theist who desperately wants not to be “alone in the universe,” it is hardly comforting that some naturalistic construal of God--say, that God is identical to humanity’s highest ideals--also gives us a *kind* of company in the universe, in the form of our ideals being with us. If anything, the linkage to the religious case reinforces, at least, the intelligibility of angsty dissatisfaction with the falsity of realism.

None of this is to say that moral realism is true; the point is that it is intelligible to regard its truth as preferable to its falsity, and that we can come to appreciate the reasons why someone might so prefer. Put another way, were we--*per impossible*--deciding in advance what kind of nature the moral domain is to have, we would have some good grounds to select from robustly realist options. There is reason to want moral justification and vindication to be wholly untethered from our responses.

---

<sup>62</sup> Nietzsche is probably the canonical debunker of this sort. Anscombe (1958) is widely regarded as offering a debunking argument along similar lines. According to the common reading, Anscombe argues that because modern moral philosophy is inextricably bound up with a Christian, legalistic conception of the world, and modern thought is non-theistic, we moderns ought to abandon modern moral philosophy for a tradition continuous with pre-Christian conceptions, especially ancient virtue ethics and an attendant moral psychology. However, because Anscombe was a devout Catholic who philosophized as such and in fact doubted the prospects of virtue ethics, this almost certainly is not the correct reading. Here I agree with Driver (2011), who reads Anscombe as offering “a *modus tollens* argument intended to establish the superiority of a religious based ethics.” See Mavrodes (1986) for a more explicit *modus tollens* in favor of a theistic ethic on the grounds of the necessity of theism for morality, in response to Mackie’s “queerness” argument.

### §2.3 *The inexpressible badness of quasi-realism*

At a high level of abstraction, anti-realist theories are theories that attempt to accommodate our moral discourse and practice without the characteristic metaphysical commitments of realism, that is, without commitments to a realistic construal of moral facts, properties, truths, and the like. The radical independence of moral truth, accommodated only by realism, can play a role in establishing an evaluative difference between realism and anti-realism partly because it amounts to a kind of independence that the standard-issue anti-realist explicitly denies, despite allowing for their own kind of degreed independence.

Things are not so easy with the quasi-realist, and this section seeks to explain (or mostly to kvetch about) why. Like the standard-issue anti-realist, the quasi-realist attempts to accommodate our moral discourse and practice without the characteristic metaphysical commitments of realism. Although the point is rarely noted, this very fact suggests that whatever goes for pro-realism *vis-à-vis* its metaphysical differences with standard-issue anti-realism should in turn apply to its differences with quasi-realism. But there is one crucial difference: the quasi-realist is an accommodationist in the extreme, seeking to, as a common boast goes, say anything that the realist can say. *Anything?* Yes—including statements that emit the strong aroma of realism, for example, “It is a response-independent, objective fact that torturing for its own sake is wrong.” Notoriously, if this project succeeds, then realists are precluded from identifying anything in the first-, second, or *n*th-order moral domain that distinguishes them from quasi-realists. And if we cannot distinguish realism from quasi-realism descriptively, it seems that we cannot say *what* distinguishes it evaluatively.<sup>63</sup> If there is no way to say how realism would be

---

<sup>63</sup> With respect to quasi-realists themselves, things did not always look so bleak. In 1973 Blackburn defined “moral realism” as the view that “the truth of a moral utterances ... consist[s] in their correspondence with some fact or

different from quasi-realism, then there is certainly no way to say why it would be better.<sup>64</sup> Due to quasi-realism's reliance on the minimalist notion of truth, according to which predicating truth of a proposition adds nothing to the assertion of the proposition itself, the problem has become known as "the problem of creeping minimalism."<sup>65</sup> The difficulty, in brief, is that if the anti-realist adopts a minimalist conception of truth that avoids identifying truth with what Crispin Wright calls "a property of intrinsic metaphysical *gravitas*,"<sup>66</sup> then all of the sorts of truth-claims that otherwise distinguish the realist from the anti-realist can no longer do their work.<sup>67</sup>

While there is not room here to delve too far into this well-trodden territory, I will raise a suspicion about the state of the dialectic between quasi-realism and realism. This will in turn shed light on what pro-realists should say about quasi-realism—or more to the point, what they should say about not being able to say much at all.

---

state of affairs", and he even declared it "false" (102). But by the 1980s, things went dramatically downhill (for the purposes of defining realism) with the increasing popularity of the minimalist theory of truth. See, for example, the struggles of Wright (1988).

<sup>64</sup> Combined with the point just made about metaphysical differences with anti-realism, we are left in a somewhat paradoxical position. Given that quasi-realism adds nothing metaphysically to anti-realism, and that realism is certainly metaphysically different from anti-realism, it seems that there *must* be some difference—we just cannot say what it is.

<sup>65</sup> See Dreier (2004) for the classic contemporary statement of the problem. Dreier laments that "It's not as if one side had better be able to come up with something clever to say about how to distinguish realism from [quasi-realism] or else the other side wins. It's rather that those of us who feel confident that there is some difference between the two meta-ethical camps should be concerned that we don't know how to say what that difference is" (31). My argument in this section says that we have reason to be *unconcerned* that we don't know how to say what the difference is.

<sup>66</sup> Wright (1995: 213).

<sup>67</sup> Many metaethicists take the problem of creeping minimalism to be a problem for quasi-realists rather than realists. The thought is that, once quasi-realists adopt the minimalist program, it is *they* who have no way of explaining why they are not full-fledged realists. If that is right, then there is little sense in a pro-realist being worried about the truth of quasi-realism, since quasi-realism is just realism. What follows in the main text can be read as pursuing the other horn of an implicit dilemma—the possibility that the problem of creeping minimalism is really a problem for realists, that it is they who have no way of explaining what they add to quasi-realism. But if one thinks that quasi-realists have no way of showing that they are not realists, then so much the better for pro-realism.

The suspicious fact is that what makes quasi-realism and realism allegedly *metaphysically* indistinguishable is a purely *semantic* maneuver—namely, the adoption of expressivism at every level of moral discourse.<sup>68</sup> But such a maneuver only works if reality is always exhausted by what we are able to say in any particular dialectical context. To illustrate, imagine two figures: a conventional moral realist and an anti-realist who is expressivist about *only* first-order moral discourse. (The early emotivists can, without too much interpretative harm, be thought of in this way.) Both the conventional realist and first-order expressivist may agree in asserting any number of first-order moral sentences, for example, “Torturing for its own sake is morally wrong.” However, they can also agree that the underlying metaphysic for each theory is different. Accordingly, they may disagree about various meta-moral, metaphysical sentences, for example, “The fact *that torturing innocent uncles for fun is wrong* is an objective, mind-independent fact in all possible worlds.” They can then point to the object of the disagreement—the nature of moral facts—and ask whether it would be better for such facts to be realistic or anti-realistic. Notice, too, that the first-order expressivist can, consistently with his own semantics, acknowledge that the kind of realistic independence highlighted in the previous section is not countenanced in his theory. He may even lament that fact.

But suppose that the first-order expressivist has an unusual conversion: he comes to endorse expressivism *at all levels of discourse about the moral*: he becomes a pan-expressivist. What becomes of the previous disagreement over the sentence, “[T]he fact *that torturing innocent uncles for fun is wrong* is an objective, mind-independent fact in all possible worlds”? Merely by changing his semantics (while doing nothing to his metaphysics), the pan-expressivist

---

<sup>68</sup> Here I am departing somewhat from the literature in identifying what we might call a *problem of creeping expressivism* rather than creeping minimalism.

now putatively agrees with the realist. Moreover, he can taunt the realist as follows: “You realists claim to be endorsing something above and beyond what we anti-realists endorse—but what?” This challenge—“but what?”—is unanswerable in the dialectical context, since there is no expressible “what” for which the pan-expressivist cannot give an expressivist reading. And so he may also taunt the pro-moral realist, “You can’t tell me what the difference is between our realism and quasi-realism, so why are you so angsty about it?”

What has the pan-expressivist *done*, dialectically speaking? I submit that he has simply taken away the realist’s license to make distinctly realistic assertions—but only in the context of confronting pan-expressivism. Unfortunately, if the realist cannot make such assertions, then the pro-realist cannot go on to say, *in light of* some expressible difference, that it is better if realism is true, even if it is! But notice, crucially, that there is no good reason to think that the limits on what can be appropriately said are likewise limits on what the world is like. Hence, the realist is pushed—coerced, really—into a narrow, involuntary mysticism. By “mysticism” I just mean any position that says there are truths that we are unable to express.<sup>69</sup> Such mysticism is *narrow*, because the pro-realist is coerced into it only when engaging a pan-expressivist. It is *involuntary*, because there are no commitments or independent considerations on the side of the pro-realist that force or even so much as motivate it. For example, there need be no explicit thought, characteristic of independently motivated mysticism, that reality outstrips our expressive capacities, though such a view is very well-motivated.<sup>70</sup> For these reasons, in the dialectical context of contrasting realism with quasi-realism, the pro-realist is coerced into going mystical

---

<sup>69</sup> This is a much thinner notion of mysticism than one finds in, for example, the canonical discussion by William James, *Varieties*, 281-283.

<sup>70</sup> For examples of mysticism that is independently motivated in this sense, see Thomas Hofweber (ms), “Are there Completely Ineffable Aspects of Reality?”; Fodor’s (1983:120ff.) discussion of “epistemic bounded[ness]”; McGinn (1989); and Chomsky (1988: 151-152).

about both the difference between realism and quasi-realism, and therefore any evaluative judgment about that difference.

While the fate of the pro-realist at the hands of the quasi-realist is a philosophical injustice, the injustice is clear enough that one need not despair. To show that the pro-realist's predicament lacks any metaphysical interest whatsoever, I will consider (and close with) three otherwise distant topics, where there are debates (either actual or merely possible) with surprisingly similar structural aspects: Berkeleyan immaterialism, theism, and the metaphysics of consciousness. In each of these debates, the pro-realist can find innocent companions in coerced mysticism, each of them falling victim to either a pan-expressivist gambit or something similar to it. Yet, in each case, it seems reasonable for the pro-realist's counterpart to remain likewise angsty, despite being robbed of expressive power. On this basis, the pro-realist may remain confident that it is better if realism rather than quasi-realism is true, even if the realist lacks the dialectical license to say it to the quasi-realist's face.

### §2.3.1 *Berkeleyan immaterialism*

According to Berkeleyan immaterialism (henceforth just immaterialism), mental substance is the only kind of substance. A common misconception has it that there are no external or physical objects of any kind according to this view, indeed no physical world at all. But immaterialism really just says something *about* such objects and the world—at the fundamental level, these are mentally constituted, and there is no need for an additional technical notion of *physical substance*. There are physical objects, but not in the kind of deep metaphysical

sense posited by *materialism*. The ordinary category “physical” is just one of many kinds of ultimately mental substance.<sup>71</sup>

Immaterialism provides an illuminating surrogate for sophisticated quasi-realist positions in metaethics. What Berkeley accommodates *vis-à-vis* materialism and dualism, quasi-realists accommodate *vis-à-vis* moral realism. Berkeley can say anything that the materialist or dualist can say. Sentences like, “There are minds and bodies,” “Some objects are physical,” “There are objects other than minds,” etc., are welcomed by the immaterialist. Yet these are just the sentences that one might have hoped would distinguish theses like substance dualism from its rivals! Crucially, Berkeley himself denies appeals to the materialist’s *technical notion* of material substance; indeed, he denies that the notion is even intelligible. Likewise, the moral quasi-realist may deny the realist’s technical notion of realistically construed, robust moral properties; indeed, quasi-realists commonly deny that such notions are even intelligible. But, with a gambit of the pan-expressivist sort, there would be no need for the immaterialist to even deny talk of “material substance,” provided it can be given an immaterialist reading.

That we are all mental substances ultimately existing as ideas in the mind of God is not an evaluatively idle idea; it is something one might appropriately want *not* to be true. Unfortunately, it is impossible to explain what is so bad about this prospect to an immaterialist who can agree with any sentences uttered by the materialist or dualist. Yet, who thinks that there *really* is no difference—both descriptively and evaluatively—between the truth of materialism and immaterialism? Those of us who are either uncommitted to either view, or committed immaterialists with a semantics that can express materialism, or committed materialists, can

---

<sup>71</sup> “[C]ollections of ideas *constitute* a stone, a tree, a book, and the like sensible things – which as they are pleasing or disagreeable excite the passions of love, hatred, joy, grief, and so forth” (Berkeley, *Of the Principles of Human Knowledge*, 1.1, emphasis mine).

express perfectly well what is lacking on immaterialism: the material world. Yet, we are simply robbed of the dialectical license to offer such a sentence in the context of arguing with an immaterialist who offers some kind of pan-immaterialist semantic program.

### §2.3.2 *Theism*

Consider a much more familiar source of angst: God. Worries about God are more likely to induce philosophical angst than immaterialism. This is not a coincidence: while the alleged badness of naturalistic religion may bear formal analogies to the alleged badness of immaterialism, religion often plays a central and existentially important a role in human life more akin to moral value.<sup>72</sup>

Although debates between traditional theists and atheists are most familiar to philosophers and non-philosophers alike, there is also a dispute between supernaturalist religion (henceforth, supernaturalism) and naturalist religion (henceforth, naturalism). Naturalists deny supernaturalist metaphysics, but seek to provide a positive account of religious discourse and practice.<sup>73</sup> Provided that something natural can play the role of the supernaturalist God, the naturalist may give naturalistic readings of various kinds of religious propositions—for example, *that God exists* or *that God loves us*. But even if the naturalist succeeds at this project of accommodating first-order religious discourse and practice, it is manifest both that naturalists and supernaturalists differ at the level of metaphysics, and that this has significant evaluative upshot. As long as the naturalist remains a second-order descriptivist, the supernaturalist and naturalist can agree that their metaphysics at the first-order level are different.

---

<sup>72</sup> Much of this, I suspect, has to do with what I say regarding “crypto-theism” above.

<sup>73</sup> See especially Johnston (2009) and Dworkin (2013).

But the naturalist may decide to play dirty. If he adopts a religious semantics in which all first-, second-, and  $n$ th-order religious discourse can come out true even without the truth of traditional theism, then the supernaturalist is robbed of his dialectical license to offer any first-, second-, or  $n$ th-order religious sentences that would otherwise distinguish the two theorists' metaphysics.

That there is no God in the traditional sense is not just an evaluatively indifferent idea; it is something one might intelligibly want *not* to be true. But it is impossible to appropriately explain what is so bad about this prospect to a naturalist who can agree with any sentences uttered by the supernaturalist. Yet, who thinks that there *really* is no difference—both descriptively and evaluatively—between the truth of supernaturalism and naturalism? Those of us who are either uncommitted to either view, or committed naturalists with a semantics that can express supernaturalism, or committed supernaturalists, can express perfectly well what is lacking on naturalism: God. Yet, we are simply robbed of the dialectical license to offer such a sentence in the context of arguing with a naturalist who offers some kind of pan-naturalist semantic program.

### §2.3.3 *Consciousness*

Finally, consider a slightly less familiar source of angst from a still-familiar topic: consciousness. Consciousness does induce philosophical angst, though perhaps not as much as naturalism or anti-realism. But it does, in many people's conceptual schemes, play an important role. For example, according to many philosophers, consciousness marks much of what is most valuable in the universe. A world of metaphysical zombies would be a much worse world than ours, at least with respect to the value contributed by consciousness itself.

Imagine a philosopher named “David.” David believes that, in addition to physical stuff in the world, there is mental stuff, in particular, there is consciousness. Moreover, David thinks that, even if we cannot quite say what consciousness is in non-mental terms, we all (provided we *are* conscious) more or less have access to the fact of its existence.

David disagrees with Phillip, who believes that there is no such thing as consciousness. Not only that, but they disagree over what they agree is a single proposition: *that there is consciousness*. It is dialectically appropriate for David to say to Phillip, “My view is different from yours, because my view posits consciousness.” This expressible descriptive difference allows them to maintain an evaluative dispute about whether it is *better* or *worse* for there to be consciousness in the world. Hence it is also dialectically appropriate for David to say to Phillip, “It would be better if my view were true, because only my view accommodates the irreplaceable, important good that is consciousness.”

But now Phillip calls over his cousin, Daniel. Helpfully for our purposes, Daniel adopts a semantics in which all first-, second-, and *n*th-order discourse about the mental can come out true even without the existence of consciousness in David’s sense. Hence, David is robbed of his dialectical license to offer any first-, second-, or *n*th-order sentences that would otherwise distinguish his metaphysics from Daniel’s. David can, of course, turn to Phillip and rightly say, “You and I both know that Daniel doesn’t *really* believe in consciousness,” but such a sentence cannot appropriately be offered to Daniel himself.

What the above three examples—material objects, God, and consciousness—show is that it is possible for an interlocutor to adopt a semantic or related program that takes away another interlocutor’s dialectical license to appropriately offer any first-, second-, ... or *n*th-order sentences in the relevant domain in order to distinguish the metaphysics of the two theories.

Consequently, it becomes impossible to give a descriptive account of what makes one *angsty* in the domain. After all, one cannot intelligibly evaluate what one cannot identify in the first place. Hence, the pro-realist is not the only sort of person coerced into what I have called involuntary mysticism. There are companions in metaphysics, philosophy of religion, philosophy of mind, and elsewhere.

What should we conclude from this, *vis-à-vis* the truth of pro-realism? Although the pro-realist is prevented from appropriately defending the claim *that it would be worse if quasi-realism were true* in the narrow dialectical context of engaging an actual quasi-realist, the assertion should be maintained in every other context. There are no good reasons for thinking that what goes for realism versus anti-realism generally, does not also go for realism versus quasi-realism. There are only good reasons to think that the pro-realist lacks the dialectical license to say so to the quasi-realist. As we have seen, the pan-expressivist gambit of the quasi-realist makes no metaphysical difference to the distinction between realism and non-quasi anti-realism; hence it must also make no evaluative difference. So, just as the realistic independence of moral vindication justifies the hope that realism rather than anti-realism is true, so it justifies the hope that realism rather than quasi-realism is true. This much, at least, we can affirm when the quasi-realist is not in the room.

### **§3 Conclusion**

In this chapter, I have commended the research program of *angsty* metaethics and the exploration of metaethical value judgments. To this end, I defended the claim that it is much better that moral realism rather than any of its rivals is true.

## Chapter 4: Preferring Moral Antirealism

Although some philosophers hope that moral realism is true, some (perhaps most) are apathetic: they do not think that anything of normative significance hinges on which metaethical theory is true—provided that one of them other than nihilism pans out. After all, normative theory is one thing, metaethics another, and it is normative theory that is responsible for telling us what is normatively significant. But a third position is neither hopeful nor apathetic. Some philosophers think that it is bad if moral realism is true, or have at least given arguments that suggest as much. In this chapter I develop a taxonomy of such perspectives and identify an underlying concern that they share. While I conclude that the arguments for hoping against moral realism are not cogent, the reasons for this conclusion are instructive for better understanding the relationship between normative and metaethics.

In brief, I distinguish between two forms of what I call *anti-moral realism*. Anti-moral realism is the position that there is something bad about moral realism in particular that does not extend to moral thought generally. Anti-moral realism divides into opposition in-principle and in-practice. According to the in-principle opposition, the truth of moral realism is intrinsically worse than its anti-realist alternatives, because something about morality itself would be worse if moral realism were true. According to the in-practice opposition, endorsement of moral realism is worse than endorsement of its anti-realist alternatives. According to this perspective, even if

moral realism is true, we should still not care about or endorse it. After showing that these arguments share a general underlying concern and structure, I respond to them.

## §1 Anti-Moral Realism

### §1.1 Killoren's argument: the unimportance of a realistic morality

The first argument that I will consider is narrowly targeted toward a particular brand of moral realism. This argument is found in David Killoren's (2016) objection to Cuneo and Shafer-Landau's (2014) version of non-naturalism. I will show how Killoren's argument can be generalized to apply to many forms of moral realism—at least, all non-naturalistic forms. Then, I will explicate another objection to moral realism and show how it shares an underlying concern and structure with the generalized form of Killoren's objection.

For the purposes of understanding Killoren's objection, the crucial thing to know about Cuneo and Shafer-Landau's view is that, according to them, moral truths obtain in virtue of conceptual truths. Consider a putative moral truth, for example, *It is morally wrong to recreationally slaughter a fellow person*. According to Cuneo and Shafer-Landau, this truth obtains in virtue of the relationships between its constitutive concepts (at least one of which, <morally wrong>, is a non-natural concept). Killoren argues that such a view precludes our having any reason to care about morality *per se*. Here is the argument:

1. Conceptual truths are normatively irrelevant.
2. According to Cuneo and Shafer-Landau, moral truths obtain in virtue of conceptual truths.
3. If X is normatively irrelevant, and Y obtains in virtue of X, then Y is normatively irrelevant. (Killoren's "Principle of Irrelevancy Transmission")  
Therefore,
4. If Cuneo and Shafer-Landau are correct, then moral truths are normatively irrelevant.<sup>74</sup>

---

<sup>74</sup> Following the argument structure of Killoren on p. 171.

I will not be concerned here to defend Cuneo and Shafer-Landau's particular brand of non-naturalism, except to make one critical comment. The Principle of Irrelevancy Transmission is fairly dubious. For one thing, it is open to Cuneo and Shafer-Landau to accept the alternative

**Principle of Relevancy Transmission:** If Y is normatively relevant, and Y obtains in virtue of X, then X is normatively relevant.

A non-naturalist who accepts the Principle of Relevancy Transmission rationally may regard it as a substantive discovery that conceptual truths are normatively relevant. The Principle of Relevancy Transmission is not obviously less plausible than the Principle of Irrelevancy Transmission, so a stalemate looms.

But I want to set such concerns aside and consider a more general--and simpler--form of Killoren's argument that does not rely on the contentious Principle of Irrelevancy Transmission. Consider the following:

1. Non-natural truths are normatively irrelevant.
2. According to non-naturalistic moral realism, moral truths are non-natural truths.  
Therefore,
3. If moral realism is true, then moral truths are normatively irrelevant.

Call this the "Irrelevancy Argument". The Irrelevancy Argument is more general than Killoren's argument, because it threatens all forms of non-naturalism that identify moral truths with non-natural truths.

More importantly for our purposes, while the Irrelevancy Argument may be deployed in defending the *falsity* of non-naturalistic moral realism, it may also be deployed in defending the *badness* of non-naturalistic moral realism (henceforth, non-naturalism). After all, it seems that it would be very bad if moral truths were normatively irrelevant. Although it is usually taken as an argument against the truth of moral realism, the result that moral truths are normatively

irrelevant is not in fact incompatible with realism. Rather than think non-naturalism is false, one might end up thinking that it is “sad but true.”<sup>75</sup>

It is easy to build on the Irrelevancy Argument and develop a somewhat more complicated argument that explicitly gives us a reason to prefer that anti-realism rather than realism is true. There are many ways to accomplish this, but here is a fairly intuitive option:

1. If moral realism is true, then moral truths are normatively irrelevant.
2. If moral anti-realism is true, then moral truths are normatively relevant.
3. Any view according to which moral truths are normatively relevant is, in that respect, preferable to a view according to which moral truths are normatively irrelevant.  
Therefore,
4. There is some reason to prefer that moral anti-realism rather than realism is true.

Call this the Relevancy Argument. The Relevancy Argument does not settle the question of which theory we should prefer, all things considered, but it offers considerable progress in answering that question in favor of anti-realism.

Notice how easy it is to defend premise 2. The Irrelevancy Argument was plausible largely because it is difficult to see how non-natural truths are normatively relevant. But it is not hard to see how natural truths are normatively relevant—at least, the sort of natural truths that appear in anti-realist theories. Here are some putative natural truths that appear in anti-realist theories:

- We are deeply opposed to the recreational slaughter of human beings.
- Idealized, fully-informed versions of our selves would disapprove of the recreational slaughter of human beings.
- Recreational slaughter of human beings is at odds with our deepest cares, commitments, plans, values, and desires.
- Recreational slaughter is an action deeply at odds with evolutionarily determined conditions in which human beings flourish.

This list is certainly not exhaustive, but it is representative of the sorts of facts that play a central role in anti-realist theories of moral truth and value. Whatever doubts one has about the truth of

---

<sup>75</sup> This is how Kate Manne (2014) describes her own internalist account of reasons, though her reasons for doing so are a bit less dramatic

anti-realist theories, it would be extraordinarily difficult to argue that they ground moral truth in something that is normatively irrelevant.

Killoren's argument, and the more generalized versions of it, are "in-principle" objections because they argue that something about morality itself is worse if realism rather than anti-realism is true. The next argument has a different structure, though I will show that the underlying concern is the same.

### §1.2 *Erdur's argument: the forbidden commitments of a realistic morality*

The second argument that I will consider is more ambitious than Killoren's original objection to Cuneo and Shafer-Landau. Melis Erdur argues that moral realism "makes the wrongness of anything always conditional on the existence of a dictate of some independent reality," and, therefore, commits realists to trying to follow these dictates whatever they may be. But, says Erdur, this commitment is itself immoral.

[S]urely, the existence of an independently issued verdict--if there were such a verdict--that genocide is wrong would not be the main or ultimate reason why it is wrong. Genocide is wrong mainly and ultimately because of the pain and suffering and loss that it involves--*regardless of whether or not the badness of such suffering and loss is confirmed by an independent reality*. The mistake in realism, therefore, is that it holds independence from our judgments in such high regard that everything else, including what really makes things right or wrong, diminishes in importance.<sup>76</sup>

This objection to realism is meant to be isomorphic to a common realist objection to anti-realism: that the anti-realist makes moral truth objectionably dependent upon our responses.<sup>77</sup>

Consider a toy version of anti-realism, say, a view according to which what is morally right and wrong is a simple function of what most people support and oppose, respectively. On this view, all it would take for lighting innocent uncles on fire for fun to become permissible is that most

---

<sup>76</sup> Erdur (2016: 597-598, emphasis original).

<sup>77</sup> For a prominent characterization of moral anti-realism in terms of the response-dependence of moral truth, see Sharon Street (2016).

people come to support such an action. It is natural to object to this by saying that, surely, the action would be wrong *even if* most people support it. Normally this is understood to be evidence that the view in question is false. But to complete the parallel to Erdur's concern, we might argue further that it is *immoral* to conditionalize one's opposition to the action on the majority opinion.

To get a better sense of the force of Erdur's objection, consider another unsophisticated version of a metaethical theory, but this time a realist one: a voluntaristic version of Divine Command Theory according to which our moral obligations are identical to whatever God happens to command, where there are no independent constraints on what these commands might be.<sup>78</sup> Divine command theorists of this sort are committed to doing horrible things if God commands them, including but not limited to lighting innocent uncles on fire for fun. It is natural to *morally* balk at this theory on the grounds that one ought not do certain things *even if God commands them*. One might think that there is something not just theoretically implausible, but morally pernicious about a theory that commits us to lighting innocent uncles on fire for fun just because God says to do so.

In Erdur's view, the same point goes for moral realism. Just as the voluntaristic divine command theorist is immorally committed to whatever behavior God happens to command, the standard-issue moral realist is immorally committed to whatever behavior independent moral reality happens to dictate. Just as you might think that one ought not do certain things even if God commands them, Erdur thinks you ought not do certain things even if they are dictated by an independent moral reality. Erdur's argument is notable partly because it issues a moral

---

<sup>78</sup> Sophisticated versions of Divine Command Theory, for example, Robert Adams' (1987) "modified" version in which moral obligations are the commands of a *loving* God, are immune from what follows.

challenge not to moral realism *per se*, but to moral realists themselves, *qua* moral agents. She intends her argument to be cogent even if moral realism is true. Here is one of several ways of putting Erdur's argument into slightly more formal terms:

1. Moral realism commits moral realists to doing whatever an independent moral reality dictates.
2. It is immoral to be committed to doing whatever an independent moral reality dictates.  
Therefore,
3. Moral realism commits moral realists to doing something immoral.

Call the "The Immorality Argument." Erdur's own conclusion is stronger than The Immorality Argument, in that she thinks one *morally ought not be a moral realist*. As far the above formulation says, there may be overriding moral reasons to be a moral realist that defeat the immorality partly constitutive of being one. But, with the exception of exotic situations in which being a moral realist is an option within a genuine moral dilemma ("Become a moral realist, or I'll light this innocent uncle on fire!"), this formulation is close enough to the conclusion that Erdur wants. Notice that, like Killoren's argument, Erdur's is consistent with moral realism being true. Just as Killoren's argument might make one think that realism is "sad but true," one might be inspired by Erdur to think the same about the immorality of being a moral realist.

Before proceeding further, it is important to remember that I am not engaged in Erdurian exegesis. Erdur herself thinks that this kind of objection applies just as much to anti-realists as it does realists. The problem is just that realists have typically thought that they were immune to an objection that they happily lobbed at anti-realists. She defends a metaphysically quietist conclusion: one ought not ground morality in anything. So, it would not be consonant with the spirit of Erdur's argument to interpret it as a reason to prefer anti-realism to realism. But Erdur's qualms about moral metaphysics in general do not prevent us from considering The Immorality Argument in isolation, as a potential reason to prefer moral anti-realism to realism.

For those who worry that there really is parity between realism and anti-realism in just the way that Erdur suggests, there is some reason to think that anti-realist commitment is at least less immoral than realist commitment. In an unpublished manuscript, “Immoral Realism,” Max Hayward argues that anti-realists are committed to abandoning morality only for *good reasons*, whereas non-naturalist realists must be committed to abandoning morality for *bad reasons*. Here is an extended excerpt from Hayward that makes the point well. His example of a moral commitment is the commitment to be faithful to one’s partner.

For a *naturalist*, we may be entirely wrong about what matters if there is no external natural world, or if we are greatly mistaken about its contents. Certain kinds of judgement dependence theorists conditionalise morality on what people want or will; if they become sceptics about other minds, they may then become moral sceptics. But these strike me as *perfectly good reasons* to suspend our commitment to morality. We all think that if there is no external world, or there are no other minds then nothing will matter morally (or what matters may be radically different). On the other hand, whether or not naturalism is true would make no causal difference to the world that we--and those we have moral relations with--live in. All our perceptual and empirical beliefs, all our inferences based on experience, could be *correct*, and naturalism still be true (or false). There might be no non-natural moral truths, and yet the rest of our beliefs about the world, and about the desires and attitudes of my partner and the commitments I have made to her, could all be correct. If I were to drop my commitment to my partner, but every aspect of our history, every aspect of her feelings and concerns stayed the same, that, I think, would constitute as deep a betrayal to her as abandoning my commitment because the Mets lost or because God is dead. ... It would be a betrayal to our fellow agents if we abandoned our commitment to them just because we decided that the natural world is all there is. Even if we never made these discoveries--never concluded that God is dead, that the Mets have lost, or that naturalism is true--simply being prepared to abandon moral commitment should that situation come to pass, or thinking that there would be nothing wrong in so doing, is itself objectionable, is itself a kind of betrayal. We ought not to conditionalise our moral commitments on any of these things.<sup>79</sup>

So, Hayward thinks there is a principled difference between (non-naturalist) realism and anti-realism *vis-à-vis* morally permissible conditionalization: realism grounds morality in something that makes no causal difference to the world, whereas anti-realism grounds morality in things that do make a causal difference. It is, in principle, permissible to conditionalize moral commitments on things that make a causal difference to the world, but it is, in principle, impermissible to conditionalize moral commitments on things that make no causal difference to

---

<sup>79</sup> Hayward (ms: pp. 13-14)

the world. Hayward's point provides a nice supplement to Erdur's argument in that it gives some reason to think that there is not the parity between realism and anti-realism that Erdur claims.

### *§1.3 Anti-moral realism and the "normative question"*

The Irrelevancy, Relevancy, and Immorality Arguments each share an underlying concern, even though they strictly speaking have different targets. The Irrelevancy Argument, recall, concludes that moral realism entails that moral truths are normatively irrelevant. The Relevancy argument concludes that, because anti-realism preserves the normative relevance of moral truths, it is in that respect preferable to realism. And the Immorality argument concludes that being a moral realist involves immorally conditionalizing one's moral beliefs on the dictates of an independent reality.

Each of these arguments relies on the idea that moral realism (and non-naturalist moral realism, in particular) provides normatively inadequate grounds for morality. Normally these arguments are marshaled in favor of the falsity of moral realism, but I have shown that they need not be. The alleged inadequacy of realist grounds has a range of deleterious results, from the unimportance of moral reality itself to the badness of basing one's moral commitments on anything that moral reality happens to say. So, the underlying idea of these arguments gives us a reason that someone might have for thinking it would be bad if realism were true, and preferring or hoping for anti-realism on that basis.

The kinds of arguments surveyed here are reminiscent of some critical remarks about realism made by Christine Korsgaard (1996). Korsgaard defines realism as the view that "moral claims are normative if they are true, and true if there are intrinsically normative entities or facts

which they correctly describe.”<sup>80</sup> Later, “Obligation is simply there, part of the nature of things. We must suppose certain actions to be obligatory in themselves if anything is.”<sup>81</sup> Korsgaard says that realism so understood “refuses to answer to the normative question,” *viz.* the question of “what *justifies* the claims that morality makes on us”.<sup>82</sup>

The thought behind Korsgaard’s critique of realism is that realists are committed to saying that we ought to obey the claims of morality simply because it is a fact that we ought to. But because it was this very fact that we were asking after in the first place, Korsgaard argues, the realist’s reply is more of a stubborn repetition than an answer to the normative question. And if realism lacks an answer to the normative question, then it turns out that realism gives us little reason to care about the demands of morality.

There is a deep connection between Korsgaard’s worry and the arguments of Killoren and Erdur that I first summarized, and then generalized, above. It is hopefully not too much of an interpretative stretch to say that Killoren and Erdur help to make concrete the objection that Korsgaard offers more abstractly. Killoren and Erdur help to show *how* and *why* realism fails to answer the normative question (or something like it) and, moreover, what is morally wrong with taking it to do so.

Recall the generalized version of Killoren’s argument (the Relevancy Argument):

1. If moral realism is true, then moral truths are normatively irrelevant.
2. If moral anti-realism is true, then moral truths are normatively relevant.
3. Any view according to which moral truths are normatively relevant is, in that respect, preferable to a view according to which moral truths are normatively irrelevant.  
Therefore,
4. There is some reason to prefer that moral anti-realism rather than realism is true.

---

<sup>80</sup> Korsgaard (1996, 19).

<sup>81</sup> *Ibid.*, 30.

<sup>82</sup> *Ibid.*, 40, 9 (emphasis original).

Premise 1 is a way of saying that moral realism fails to answer Korsgaard's normative question. After all, if realism *did* answer the normative question, then it could not be that moral truths are normatively irrelevant. If this is right, then premise 3 can be read as saying that metaethical theories that answer the normative question are, in that respect, preferable to theories that don't.

Now recall the formalization of Erdur's argument:

1. Moral realism commits moral realists to doing whatever an independent moral reality dictates.
2. It is immoral to be committed to doing whatever an independent moral reality dictates.  
Therefore,
3. Moral realism commits moral realists to doing something immoral.

This argument is more distant from Korsgaard's concerns, but it still bears a significant and illuminating connection to it. Suppose we ask *why* it is immoral to be committed to doing whatever an independent morality dictates. The answer must be that there is no good reason to grant an independent reality the relevant kind of *authority* over one's actions. But this is just a way of saying that positing an independent moral reality does nothing to answer the normative question. If independent moral reality were a sufficient basis for answering the normative question, then it would no longer seem immoral to conditionalize one's moral commitments upon its dictates. If these connections are right, then Erdur's argument is a way of drawing something out of Korsgaard's concerns that is of distinctly moral significance to metaethicists themselves: it is immoral to adhere to a theory that fails to answer the normative question.

## §2 Responding to anti-moral realism

I am what I call a pro-realist. That is to say, I think it would be better if moral realism were true, rather than its alternatives. So, I have an interest in responding to the above suite of arguments. To that end, I will argue that each of the above arguments elides the distinction between normative and metaethics, albeit in subtle ways. When this distinction is clarified and held in view, the arguments lose their force.

## §2.1 Killoren

First, Killoren. Recall Killoren's Irrelevancy Argument, central to which was the Principle of Irrelevancy Transmission:

**Principle of Irrelevancy Transmission:** If X is normatively irrelevant, and Y obtains in virtue of X, then Y is normatively irrelevant.

Killoren motivates the Principle of Irrelevancy Transmission itself partly by appealing to the example of cultural relativism. He points out that the cultural relativist may give the following unconvincing reason for why recreational slaughter of persons is wrong: "Our society is such that recreational slaughter of a fellow person is condemned by our society's generally accepted moral standards" (171). Yet, says, Killoren, "the mere fact that our society happens to generally accept a standard that forbids a given action does not seem, on its own, to provide any reason to refrain from performing that action." Hence, the relativist's reason fails to explain why recreational slaughter is really wrong, and in fact undermines that claim. This failure is itself evidence for the Principle of Irrelevancy Transmission, because the principle nicely *explains* why the normative irrelevance of what our society happens to accept in turn seems to undermine the claim that we have reason not to engage in recreational slaughter.

While I concede that this example makes *some* version of the principle plausible, at the same time it casts doubt on whether the plausible version is relevant to *metaethics*. The example makes salient something fishy about the relationship between the principle and grounding projects in metaethics more generally. As described by Killoren, cultural relativism is a position in *normative* moral theory, in which grounding projects seek precisely to identify the general normative grounds for true moral propositions. But grounding projects in metaethics are engaged in something different from grounding projects in metaethics. Some normative fact, N,

may be normatively grounded in some further fact, X, yet be metaphysically grounded in another fact, Y. Consider the following claim that, let us suppose, is true right now: *you ought to clean your room*. This truth may be grounded in the fact that *your parents commanded you to clean your room*, and that fact may be normatively relevant in virtue of some underlying fact about what is owed to parents. But none of this involves any metaethics—in particular, it does not involve any moral metaphysics. After all, imagine that someone asks for a justification for why they ought to clean their room. It is sensible to respond, “Because your parents told you to,” but quite ridiculous to respond, “Because it’s an irreducibly normative, non-natural conceptual truth that you ought to obey your parents.” The lesson here is that normative justification is not a species of metaphysical ground, and the Principle of Irrelevancy Transmission applies to the former and not the latter.

Once we make the distinction between normative and metaphysical grounds, neither the Irrelevancy nor Relevancy Arguments are plausible. This is because all of them rely on the thought that the metaphysical grounds of moral truths are normatively irrelevant, whereas they would need to show that the normative grounds are normatively irrelevant. Concluding that morality is normatively irrelevant because its metaphysical grounds are normatively irrelevant is a bit like concluding that your Grandfather’s watch is insignificant because the metal substance in which it is grounded is insignificant. In fact, the watch’s significance is not determined by its metaphysical grounds.

## §2.2 *Erdur*

Let us turn now to Erdur’s argument. Recall that Erdur thinks that moral realists are committed to conditionalizing their moral views on the dictates of an independent reality. This is because she thinks that a person’s metaethical view is the correct place to terminate a chain of

normative why-questions. To illustrate, consider the following exchange between Eleazar (a layperson) and Esther (a Kantian), which I have loosely adapted from pp. 593-594 of Erdur's paper.

<b>Esther:</b>	A's killing of B was not wrong.
<b>Eleazar:</b>	Why?
<b>Esther:</b>	Because it was self-defense.
<b>Eleazar:</b>	Why does the fact that A's killing was self-defense count against its wrongness?
<b>Esther:</b>	The fact that A's killing of B was self-defense makes a moral difference because it makes a difference to the universalizability of A's maxim--killing to save one's own life passes the universalization test, whereas killing for another reason might have failed it.

Suppose now that Eleazar is still unsatisfied. How might the exchange naturally proceed? Erdur imagines something like the following, in the case that Eleazar is not only a Kantian but a moral realist:

<b>Eleazar:</b>	Why does the universalizability of a maxim render an action permissible?
<b>Esther:</b>	Because an independent moral reality dictates that it does.

Once a "why" question is asked about an abstract normative theory, the appropriate next step, according to Erdur, is to ascend to the level of metaethics. Metaethical theories, therefore, "may naturally be heard as very general substantive moral claims about why (in the end) right things are right and wrong things are wrong" (595).

Given what I have said in response to Killoren, my response to Erdur should be unsurprising. There is something confused about answering a substantively moral why-question with a metaethical statement--and, especially, a *metaphysical* metaethical statement. Philosophical taxonomy itself casts doubt on metaethical answers to normative moral questions. After all, as I argued in the previous subsection, normative justification is one thing, metaphysical grounding another.

To see the broader point more concretely, it is useful to depart from moral philosophy for the moment. Consider the following exchange about physics between Newton and Moses, intended as a scientific parallel to the exchange imagined by Erdur:

<b>Newton:</b>	X's state of being at rest hasn't changed.
<b>Moses:</b>	Why?
<b>Newton:</b>	Because it has not been impressed upon by any force.
<b>Moses:</b>	Why does the fact that an object has not been impressed upon by any force make a different to its state of being at rest?
<b>Newton:</b>	The fact that an object has not been impressed upon by any force makes a different to its state of rest because of the Law of Inertia--every body persists in its state of being at rest or of moving uniformly straight forward, except insofar as it is compelled to change its state by force impressed.

So far, so good. But suppose that Moses is still unsatisfied. How might the exchange naturally proceed? Here is what the analogue to Erdur's ascent to metaethics would look like, if we think of Newton as a scientific realist:

<b>Moses:</b>	Why does the Law of Inertia make any difference to whether an object is at rest?
<b>Newton:</b>	Because an independent reality dictates that it does.

It is immediately recognizable that Newton's final answer is not informative *as a piece of physics*. If Newton is doing physics, then his final answer does not even so much as add anything to what he has said previously. As a sentence in physics, the statement that an independent reality dictates that the Law of Inertia makes a difference to whether an object is at rest is just a restatement that the Law of Inertia makes a different to whether an object is at rest.

But Newton's final answer is, or would be, informative as a piece of *philosophy of science* or, in particular, the *metaphysics of physics*. In such a context, we would think of his answer as a statement of scientific realism. (I am here conceiving of philosophy of science in relation to scientific discourse and practice as analogous in relevant ways to metaethics in relation to moral discourse and practice.) Unfortunately, in the present context, informativeness comes at the expense of relevance. Moses just wasn't asking about whether scientific realism is true.

All that being said, Newton's final answer is relevant to other questions that Moses could have asked. Here is just one such question that is better-suited to Newton's final answer:

**Moses:** When you say that the Law of Inertia makes a difference to whether an object is at rest, are you saying that there *really* is such a law, written into the fabric of the Universe, or is that just a construct that helps us carry on scientifically?

Newton's final answer is relevant to this question, because this question is located squarely in the domain of philosophy of science, and the metaphysics of physics in particular.

The same analysis applies to the exchange between Esther and Eleazar. Esther's final appeal to "an independent moral reality" adds nothing to the discussion if it is an answer to a question within normative ethics. Yet, "Why does the universalizability of a maxim render an action permissible?" is a question in normative ethics, and so Esther's final answer about the dictates of an independent reality is either uninformative or irrelevant. If it is an answer in normative ethics, then it simply repeats what came before. If it is an answer in metaphysics, then it is irrelevant to Eleazar's normative question. But we can change Eleazar's question just as we changed Moses' question. Here is a question better-suited to Esther's final answer:

**Eleazar:** When you say that the universalizability of a maxim renders an action permissible, are you saying that there *really* is such a moral law, written into the fabric of the Universe, or is that just a construct that helps us carry on morally?

Esther's final answer is relevant to this question, because this question is located squarely in the domain of metaethics.

Once we disambiguate normative from metaethical "why" questions (and their "because" answers), we should not accept Erdur's claim that metaethics provides the termination point for chains of normative why-questions. And if we do not accept Erdur's claim that metaethics provides the termination point for chains of normative why-questions, then we need not worry

about normative commitments being conditional on metaethical positions. Since this claim about commitment is essential the Immorality Argument, the argument fails.

Before finishing this subsection, it is worth applying its lesson more concretely. According to Erdur, moral realists are committed to denying the following proposition (among infinitely many others): *Genocide is wrong mainly and ultimately because of the pain and suffering and loss that it involves--regardless of whether or not the badness of such suffering and loss is confirmed by an independent reality*. I have shown that there is an equivocation between the normative and metaethical “because.” As grounds for wrongness of genocide, *pain and suffering* are not competitors with *the dictate of an independent reality*. Rather, the realist view is that the ultimate normative grounds for the wrongness of genocide are to be understood metaphysically as constituted by an independent reality. It is not that the independent reality issues a *further* dictate beyond the proposition, *Genocide is wrong because of the pain and suffering and loss that it involves*. Rather, *that* dictate just *is* an independent reality.

Erdur imagines that realists (and antirealists) have the second-order committed to give up their first-order moral commitments in the event that their metaphysics tells them to do so. Once we have rightly separated the roles of normative and metaethics, it is clear that this hypothetical is not plausible. Using Erdur’s terminology, the only “reality” that can tell us what to do is normative reality. Hence, as far as theories go, it is the best normative moral theory that should tell us what we morally ought to do. The best metaethical theory tells us the metaphysical, epistemology, semantic, and perhaps psychological status or nature of the best normative moral theory.

If this is the right way to understand normative and metaethics, then the closest hypotheticals to Erdur’s will not--indeed, cannot--have the morally objectionable results that

she worries about. The *only* sense in which an “independent moral reality” can “dictate” something—say, *that p is morally wrong*—that opposes what we currently think is by constituting the metaphysical, epistemological, semantic, and psychological underpinnings of *the true normative theory’s* dictation *that p is morally wrong*. But it is hardly morally objectionable to be committed to thinking *that p is morally wrong* in the event *that p is morally wrong*.

### §2.3 Korsgaard’s normative question—and Parfit’s principled non-answer

What goes for Killoren and Erdur goes also for Korsgaard’s normative question. We should not expect that metaethics will give us a *normative reason to care about morality*. Even though I do not think that metaethics is evaluatively or normatively idle, I do hold that it is idle in *this* particular sense.<sup>83</sup>

Recall that Korsgaard thinks that the realist appeal to mere facts is insufficient to answer the normative question. This is because, “If it is just a fact that a certain action would be good, a fact that you might or might not apply to deliberation, then it seems to be an open question whether you should apply it” (317). In response, Derek Parfit found himself mystified that anyone could endorse a proposition like *It is a fact that I ought to X* and simultaneously wonder *Why should I X?* In response to Korsgaard, Parfit writes, “If it is a fact that you should do something, it is not an open question whether you should do it” (418). Parfit felt that, once you accept the fact that you ought to X, there is no further question for you about whether to X. Philosophers who claim that non-natural facts about what one ought to do are not necessarily normative struck Parfit as conceptually confused. I confess to deep sympathy with Parfit’s reaction to Korsgaard’s claim. If we know that it is a fact that I ought to X, then we know that the

---

<sup>83</sup> In Chapter 2, I gave arguments for why and how we should expect that metaethical theories have significant—and different—evaluative upshot.

true normative theory tells me to X. The *normative* grounds for why the true normative theory tells me to X may be open in an epistemic sense—for example, I may not know what they are—but this is not a failing that can be pinned on the realist.

In explicating this point, it helps to emphasize the difference between the thing in virtue of which a claim is true and the thing in virtue of which that same claim is normative. The proposition *recreational slaughter of persons is wrong* may be *true* in virtue of its constitutive concepts, but that need not be why it is normative. It is normative, we might plausibly argue, in virtue of the unjustified harms that it causes. This distinction, between something's being true and its being normative, helps to explain the odd sound of simultaneously judging that X is wrong and that it is normatively insignificant that X is wrong.

### §3     **Are normative and metaethics normatively unrelated?**

Based on my replies to Killoren, Erdur, and Korsgaard, one might suspect that I want to isolate metaethical theorizing from normative theorizing entirely. Moreover, it might seem that I think that anti-moral realism itself necessarily involves a confusion. But this is not the case. In my view, metaethical theories are just as susceptible to normative (including moral) evaluation as other kinds of philosophical theories. Consider the example of theism in the philosophy of religion. Theism is not a substantive moral theory, or a moral theory of any kind, but this does not make theism immune from normative or moral critique. Recently, in fact, philosophers have devoted considerable attention to questions about whether (and in what ways) it would be good or bad if theism were true.<sup>84</sup> We can easily explore these questions without conflating normative ethics and philosophy of religion; the fact that we can wonder whether or not God's existence would be good does not mean that theism is a normative theory.

---

<sup>84</sup> See especially the volume edited by Kraay (2018).

Although these questions are seldom pursued, we can ask the same sorts of normative or evaluative questions about metaethics, without needing to think that metaethics is really just normative ethics in disguise. We can ask questions like, “Would it be better if moral realism rather than antirealism were true?”, or “Would it really be so bad if nihilism were true?”<sup>85</sup>, or any number of other similar normative questions *about* metaethics.

So, my complaint about eliding the distinction between normative and metaethics is a complaint about equating the two domains, but *not* about relating them at all. Questions about the normative or evaluative upshot of different metaethical theories are important and interesting questions that philosophers should devote more attention to, even if the arguments due to Killoren, Erdur, and Korsgaard somewhat miss the mark.

#### **§4 Conclusion**

I have addressed several arguments for preferring anti-realism to realism. One argument has to do with what realism says about morality itself. According to this argument (an in-principle objection), morality itself would be worse if moral realism were true. Another has to do with the badness of endorsing realism (an in-practice objection). According to this second argument, the realist is necessarily committed to wrongly conditionalizing first-order moral judgments on metaethical posits. I have shown that these arguments exhibit a shared structure and underlying concern about the normative insignificance of realist metaphysics, one that resembles a concern that was expressed in general form by Christine Korsgaard (1996). Finally, I have argued that the underlying concern can be answered by keeping in view the distinction between normative and metaethics. In so doing, I have not argued one way or the other about

---

<sup>85</sup> Guy Kahane recently asks this question in his, “If Nothing Matters.” I ask both questions in Chapter 3.

whether it really is good or bad if moral realism or anti-realism is true. That is a task for other papers. I have merely undercut some interesting and powerful reasons for favoring anti-realism.

## Chapter 5: Value as a Guide to Truth in Metaethics (and Other Domains)

Philosophers sometimes argue that certain views entail bad consequences. For example, Derek Parfit argued that subjectivism about reasons entails that nothing matters.<sup>86</sup> Robert Adams argued that modal realism would justify moral indifference.<sup>87</sup> Several philosophers and non-philosophers alike argue that atheism entails that life is meaningless. Intuitively, these arguments have little to do with what is actually true about reasons or modality or God and speak, at most, to what we should want or hope for. We resist the inference from badness to falsity (or from goodness to truth) in these cases for the same reason that we resist it anywhere else: it seems to be a bad inference. Call this seemingly bad kind of inference a “value-truth inference.” By this I mean an inference from the evaluative upshot of *p* (whether *p* would be good or bad) to the semantic properties of *p* (whether *p* is true or false). In this chapter I cast some doubt on the traditional reasons for rejecting value-truth inferences and defend a positive account according to which value judgments in some contexts of inquiry may provide some evidence of truth. In brief, the contexts of inquiry in which value can be evidence of truth are those in which theorists are at least partly responsible for accommodating the distinctively evaluative features of the objects of their inquiry. Finally, I argue that my framework is superior (both in its generality and

---

<sup>86</sup> Parfit (2011).

<sup>87</sup> Adams (1979).

relative simplicity) to some recent approaches to this and related questions developed in normative theory.

## §1 Traditional objections rejected

There are a number of attractive explanations for why value-truth inferences seem inappropriate. One substantive reason has both an empirical and logical gloss. The empirical gloss is just that there are many horrible truths—for example, that many human beings lack adequate food, or that wars occur. The logical gloss is that there is no logical incompatibility between something's being both horrible and true.<sup>88</sup> Consider that we do not generally think that something is true *despite* its being horrible; we just think that something can be both true and horrible. Put simply, the mere fact that something would be better seems like the wrong kind of reason for believing that it is true. That something would be better is evidence of goodness, but not of truth.

Though attractive, these reasons are not sufficient to defeat value-truth inferences. Both the nature of value-truth inferences, as well as the reasons that supposedly contravene them, are underspecified. To illustrate, consider the following informal inference: *A loving God is better than an unloving God, so God is loving*. On the surface, this is a value-truth inference in the relevant sense. The premise, *A loving God is better than an unloving God*, is a judgment of comparative value, and the conclusion, *God is loving*, is a judgment of non-evaluative truth. Nevertheless, it is a mistake to reject this inference on the grounds of the considerations stated previously, for it is pretty clearly enthymematic. Here is a charitable interpretation of the whole inference:

---

<sup>88</sup> Cf. Kasper Lippert-Rasmussen on rejecting value-truth inferences in discussions of inviolable rights: "Someone who says 'I agree that it would be better if we were inviolable, but I think we are not' does not contradict himself" (136).

- [A1] God is the greatest possible being.
- [A2] A loving God is better than an unloving God.  
Therefore,
- [C1] God is loving.

A1 says something specific about God in virtue of which value judgments about God have implications for the truth about what God is like. Filling out the argument in this way, do we still have a value-truth inference, in the sense this paper opened with? If we do, then we have an example that is in tension with the considerations cited above in opposition to such inferences. In the case of God's nature, it seems that there *is* an incompatibility (or at least a tension) between something's being both horrible and true. Likewise, something's being better--again, at least in the case of God's nature--does not seem like the "wrong kind of reason" for believing that it is true.

One possible way of distinguishing between the God case and problematic value-truth inferences is to mark out the latter as distinctively lacking the sort of linking premises available in the God case. Perhaps the God case has what we might call a "linking premise," that is, a premise that, by *fiat*, simply makes value relevant to truth. To see the point, consider hilarity-truth inferences, which are inferences from something's being hilarious to its being true. When we reflect abstractly, these seem like bad inferences. After all, there are plenty of unfunny truths, and there does not seem to be any logical incompatibility between something's being both unfunny and true. Put simply, that something would be hilarious seems like the wrong kind of reason for believing that it is true. But now imagine that someone offers the following counterexample in defense of hilarity-truth inferences:

- [A3] God is the most hilarious being.
- [A4] A loving God is more hilarious than an unloving God.  
Therefore,
- [C2] God is loving.

Like A1, A3, by a kind of *fiat*, neutralizes the problem with inferring from a seemingly truth-irrelevant property to truth, by linking the two together.

Perhaps the genuinely problematic examples of value-truth inferences are just those inferences from value to truth that lack property-truth linking premises like A1 and A3. But now, one begins to lose a sense of both the interest of the inferences in question, as well as the non-triviality of the objections to them. After all, it seems that for *any* object, *x*, and *any* logically and metaphysically independent substantive properties, *F* and *G*, the mere fact *that x is more G if it is also F* is not a reason to think that *x is F*. In other words, inferences patterned after the following schema are generally invalid without linking premises:

- [A5]    If *x* is *F*, then *x* is more *G* than if *x* is not *F*.  
          Therefore,  
[C3]    *x* is *F*.

This inference requires an additional premise of the following form:

- [A6]    *x* is maximally *G*.

Notice that there need not be anything essentially normative or evaluative about *G* (or *F*) for this kind of inference to require supplementation. At most, what is needed is that *G* admits of degrees, such that it is intelligible to say that *X* is more or less, or minimally or maximally, *G*.

The arguments and observations of this section put considerable pressure on the idea that value-truth inferences run afoul of the existence of horrible truths, or that a value judgment provides the wrong kind of reason for a judgment of truth. In fact, the attribution of any property, by itself, in the context of the argument schema above, provides the wrong kind of reason for a truth judgment—without additional linking premises. So, the pertinent question about value-truth inferences seems to be whether (and if so, where) there are plausible value-truth linking premises that can both validate them and explain their (occasional) intuitive appeal.

## §2 Some exotic value-truth linking premises

The substantive debate about value-truth inferences ought to be about which value-truth linking premises are independently plausible. In the case of God, the overwhelming judgment of philosophers who have written on these matters is that A1 (*that God is the greatest possible being*) is much more plausible than something like A3 (*that God is the most hilarious being*). I suspect that what makes philosophers suspicious of value-truth inferences in general is the thought that, as far as both general empirical and metaphysical matters are concerned, there are typically no plausible linking premises on offer.

Leaving God to one side, we can emphasize both main points—that value judgments are not the wrong kinds of reasons *per se*, and that the proper ground of suspicion lies elsewhere—by mentioning two other perspectives that, if they are true, make many value-truth inferences more plausible. Consider axiarchism. Axiarchism is a kind of Platonism, defended especially by John Leslie, according to which the world exists because and for the sake of goodness.<sup>89</sup> The good is both the world's cause and its purpose—that is, the world both derives from and aims at the good. Thomas Nagel sketches a version of this view in *Mind and Cosmos*.<sup>90</sup> According to Nagel, value (along with mind) is a fundamental constituent of the cosmos and something that cosmic history, in some sense, aims to produce. If you accept either of these frameworks or something similar, then the fact that some state of affairs is good, or would be better than some other state of affairs, may provide some reason to think that that state of affairs either has or will come about.

---

<sup>89</sup> John Leslie (1979). See also Derek Parfit (2004). See Tim Mulgan (2017) for discussion of this family of views.

<sup>90</sup> Thomas Nagel (2012).

Second, consider the less metaphysical (but no less bold) *moral faith* illustrated by the famous declaration, “The arc of the moral universe is long, but it bends toward justice.”<sup>91</sup> One might take this attitude toward history without being an axiarchist like Leslie, Parfit, or Nagel. From this point of view, the fact that some state of affairs, say, a more just society, would be better, is in fact a reason to believe that it will come about (and, hence, is available as a reason for action).<sup>92</sup>

These perspectives--axiarchism, on the one hand, and a kind of moral or political faith, on the other--are quite radical, and most philosophers are bound to find them implausible. But perspectives such as these are nevertheless candidates for providing the kind of background information that could validate inferences from value to truth. These theories can provide value-truth linking premises that license inferences from value judgments to truth judgments. In the next section, I will offer a theory about what can (much more plausibly) provide value-truth linking premises, using metaethics as my example domain. The chief advantage of my account over the two perspectives addressed in this section is that I need not posit anything metaphysically exotic about the world as a whole.

---

<sup>91</sup> The exact quote was made famous in a speech by Martin Luther King, Jr. in 1964, but it originates from Theodore Parker, “Of Justice and Conscience,” in *Ten Sermons of Religion* (Boston: Crosby Nichols and Company, 1853). In Parker’s words, “Look at the facts of the world. You see a continual and progressive triumph of the right. I do not pretend to understand the moral universe, the arc is a long one, my eye reaches but little ways. I cannot calculate the curve and complete the figure by the experience of sight; I can divine it by conscience. But from what I see I am sure it bends toward justice” (pp. 84-85). See also Preston-Roedder, “Faith in Humanity,” for a related account of faith, particularly in persons.

<sup>92</sup> There is an interesting connection here to the literature on moral explanations. If, for example, Sayre-McCord and Sturgeon are right that moral facts can help to explain events that *have* happened, then there is no immediately obvious barrier to their helping to predict events that *will* happen. The role of moral facts in prediction is underexplored.

### §3 Value as a guide to truth in metaethics

Let us turn now to the case of metaethics, a domain in which some philosophers make strong value judgments about what would be better (see the reference to Parfit at the beginning of this paper), but generally do not think such judgments are relevant to truth. As Kate Manne suggests of her own internalist view of reasons, some metaethical theses may be “sad but true.”<sup>93</sup> If I am right in my account of what value-truth inferences are, and why it is that philosophers are skeptical about them, then the right question to ask is not whether it is appropriate to infer a putative metaethical truth from its being good, but whether there are any plausible linking premises that will allow us to move from value judgments about morality to the correct account of morality.

To begin, here is an example of a zealous value-truth inference in metaethics:

- [A7] Moral norms are maximally authoritative.
- [A8] Moral norms are more authoritative if they are objective than if they are subjective.  
Therefore,
- [C4] Moral norms are objective.

Both A7 and A8 are controversial, but the argument from A7-C4 exhibits the right structure for inferences from value judgments about ways morality might be, to truth judgments about the way that morality is. Here is another zealous argument:

- [A9] The moral domain is maximally worth caring about.
- [A10] A moral domain, construed realistically, is more worth caring about than the moral domain, construed antirealistically.  
Therefore,
- [C5] The moral domain is realistic.

Both A7-C4 and A9-C5 are zealous—but are they overzealous? They both have a particular species of linking premise that I will call “maximalist premises”, which strongly resemble A1 in

---

<sup>93</sup> Kate Manne (2014). According to Manne, internalism about reasons is sad, even “gloomy”, because it precludes our ability to take the otherwise appealing moral stand that certain particularly evil people (who are deeply and consistently devoted to doing evil) have a reason to do better.

the earlier God case. A maximalist premise says that something has some property to the maximal degree. In the case of morality, such premises call to mind doctrines about the overriding authority of moral demands. I happen to think that some kind of overridingness thesis is true, if out of fashion. But in recognition of the contemporary unpopularity of the thesis, it is useful to know if there are inferences in the neighborhood that do not make use of controversial maximalist premises.

It seems to me that maximalist premises are especially useful for deductive value-truth inferences. They are a handy way to get from a comparative claim (that some way of being is better than another) to an absolute claim (that something really is the better way). After all, anything less specific would be too underdetermined. That the moral domain is worth caring about, or even very worth caring about, does not, even in conjunction with A10, tell us anything about whether realism or antirealism is true. That is because, as far as anything our replacement premise plus A10 says, we have no idea to what degree morality is worth caring about, either in itself or according to various metaethical views. For all the replacement premises say, moral realism might make morality *too* worth caring about!<sup>94</sup>

Fortunately, value-truth inferences need not be deductive. We may also make inductive or abductive value-truth inferences, which are inferences from something's being better a certain way to its probably being that way, or best explained as being that way. Or, even more weakly, we may infer from something's being better a certain way to its being, just to that extent, probably that way. Here is an example adapted from A9-C5:

---

<sup>94</sup> This kind of worry would resemble so-called "demandingness objections" in normative moral theory.

- [A11] Morality is very worth caring about.  
[A12] Moral realism better explains A15 than moral antirealism.  
Therefore,  
[C6] Moral realism is more probable than antirealism in at least one respect.

A11-C6 is a value-truth inference, because it involves moving from a value judgment (A11, that morality is very worth caring about) to a truth judgment (C6, that moral realism is more probable than antirealism in at least one respect). I regard C6 as a truth judgment in the relevant sense, because probability is being used in a purely alethic way. (A judgment can be normative without being evaluative, in the sense relevant here.) To say that moral realism is made more probable than anti-realism by some phenomenon is to say that, *at least with respect to that phenomenon*, it is more likely that moral realism *is true*.

Unfortunately for the iconoclastic defender of value-truth inferences *per se*, but fortunately for the somewhat deflationary purposes of this paper, our best candidates for such inferences begin to look suspiciously like more ordinary inferences made in metaethical—and every other kind of—descriptive theorizing. All A11-C6 is really saying is that morality has some property, F, and that moral realism better accommodates F than moral antirealism. But this is just an ordinary way of conducting philosophical theorizing in general. Compare: many philosophers have argued that moral realism better accommodates the surface-level grammar of moral sentences than moral antirealism, and so this property speaks in favor of the truth of realism. The surface-level grammar of moral sentences does not seem like an evaluative property of morality, yet arguments for moral realism on their basis seem structurally identical to A11-C6. In metaethical value-truth inferences, it just so happens that the property is evaluative in nature.

Is there anything special, then, about value-truth inferences in general, or in the particular moral case? I submit that there is nothing special about these inferences, and

consequently that there should also be nothing, on reflection, suspicious about them. The reason why value-truth inferences are attractive and appropriate in the metaethical case, as in the theological case discussed in the previous section, is simply because theorists in these domains are responsible for accommodating both evaluative and non-evaluative aspects of the target phenomenon. In some of the most paradigmatic domains of rational inquiry, e.g., all domains of natural science, theorists are responsible for accommodating only non-evaluative phenomena. Physicists, for instance, must accommodate what we know from other natural sciences, what we observe in the laboratory, various theoretical *desiderata*, and so on, but they need not accommodate distinctly evaluative phenomena like the fact that persons are valuable, or that the world is valuable, or even that physics itself is valuable. Likewise, evolutionary biologists, even though they have plenty to explain about the nature of life, are not responsible for accommodating evaluative facts, for example, that it is wrong to end a life for no reason.

In the cases of both theology and moral theory, the objects of theorizing (God and morality, respectively) have evaluative properties that we ordinarily want theorists to accommodate. A good theologian not only accounts for, say, God's having created the world, but also accounts for God's goodness. If a theology fails to fully accommodate God's goodness, then that is a theoretical cost for the theory--not merely the positing of a horrible truth, though it may also be that. Likewise, a good metaethicist does not only account for, say, the surface-level grammar of moral sentences, but also for morality's importance or value. If a metaethical theory fails to fully accommodate morality's importance, then that is a theoretical cost for the theory--not merely the positing of a horrible truth, though it may also be that.

The examples of metaethics and theology are sufficient to show that what validates value-truth inferences has to do with the responsibilities bestowed on theory-builders by their

respective domains. But it is worth wondering whether there are cases other than metaethics and theology where philosophical theorizing is beholden to evaluative aspects of the target phenomenon.<sup>95</sup> There are. For example, insofar as knowledge is thought to be valuable or worth having, accounts of knowledge are responsible for accommodating this fact.<sup>96</sup> There are other cases of philosophical theorizing that less obviously resemble value-truth inferences, but which have exactly the motivations laid out here. For example, in his influential paper, “Freedom of the Will and the Concept of a Person,” Harry Frankfurt seems to presuppose that theories of personhood are better if their truth means that personhood is itself better. He writes,

In those senses of [“person”] which are of greater philosophical interest, ... the criteria for being a person do not serve primarily to distinguish the members of our own species from the members of other species. Rather, they are designed to capture those attributes which are the subject of our most humane concern with ourselves and the source of what we regard as the most important and most problematical in our lives.<sup>97</sup>

At the end of the paper, Frankfurt claims theoretical victory for his account of freedom of the will partly by claiming that his account “satisfies another condition that must be met by any such theory,” viz., that “freedom of the will should be regarded as desirable.”<sup>98</sup> It is not clear why it should count against a theory of freedom that freedom comes out as being not as desirable as we

---

<sup>95</sup> Some philosophers with whom I have discussed these ideas suggest that, contrary to what I have said so far, the natural sciences do involve accommodating evaluative features of the domain, because theoretical *desiderata* (mentioned above) are normative. This misses my point, but in an illuminating way: theoretical *desiderata* are not features of *the target domain of inquiry*. Simplicity, for example, is thought by many to be a property of good theories and might be, in that sense, normative for theorizing itself. But simplicity is not presumed in advance to be a property of, say, biological organisms or planets—certainly not in any normative way. Of course, theorists working on the nature of theorizing itself are partly responsible for accommodating the evaluative aspects of the theoretical virtues. But *theorizing* is a very different kind of domain from the domains of the natural sciences themselves.

<sup>96</sup> Consider Edward Craig’s *Knowledge and the State of Nature*, which connects knowledge to the importance of having good informants about the world. Craig’s theory gains a theoretical advantage insofar as it accommodates what is *valuable* about knowledge. A theory that renders knowledge valueless is not for that reason necessarily false, but it can for that reason only be true at a cost.

<sup>97</sup> Frankfurt (1971, 6).

<sup>98</sup> Frankfurt (1971, 17).

thought it was, unless we take theorists in this domain to be responsible for accommodating certain evaluative as well as non-evaluative aspects of freedom itself.

Before showing the advantages of my account over its rivals, let me recap what has been established. First, the traditional reasons for rejecting value-truth inferences, though intuitively attractive, are not very good. Once we see why they are not very good, we can also see that what bothers us about value-truth inferences could be a problem with any other kind of property, evaluative or not: the absence of a property-truth linking premise. Finally, one need not resort to metaphysical esoterica like axiarchism in order to vindicate value-truth inferences. Rather, I have defended a self-consciously boring account: value judgments can provide some defeasible evidence for the truth of a theory about some domain just when theorists of that domain are responsible for accommodating certain of its evaluative features. The proposal is *boring*, because there is nothing special about evaluative features as far as inferences to truth are concerned. In principle, the point applies to any other kind of feature. If theorists were responsible for accommodating the comedic features of a domain, then the fact that some theory secures more hilarity than another would be *some defeasible evidence* for that theory.

#### **§4 Comparative advantages of the view**

Due to their appeal in first-order moral theorizing, the recent literature on value-truth inferences focuses somewhat narrowly on moral philosophy. Generally speaking, extant theories either suffer from being applicable only to moral theory, or are needlessly complicated or esoteric (much like axiarchism). The point in this section is decidedly not that there is nothing true or insightful in the approaches discussed in what follows: in fact, I think that each of them shows us something important about morality and normativity. Rather, whatever these theories have to

say, I aim to show that they are not needed for the specific purpose of defending value-truth inferences.

#### 4.1 Nagel's "curious type of argument"

First, a word on how the contemporary literature got started. Thomas Nagel is largely responsible for the recent literature on value-truth inferences as they appear in normative moral theory, though he primarily notes the appeal of such inferences in moral theory rather than attempts to explain them. In "Personal Rights and Public Space," Nagel defends and employs a broadly deontological approach to human rights, according to which rights are intrinsic to persons rather than instrumental. Nagel argues that, because it is better if we have these rights, we--therefore--do in fact have them. In his words,

The argument is that the most plausible alternative morality, which is based solely on the agent-neutral value or disvalue of the actual enjoyment or infringement of certain freedoms, and so on, fails to give any place to another very important value--the intrinsic value of inviolability itself. The argument is that we would all be worse off if there were no rights--even if we suffered the transgressions which in that case would not count as violations of our rights--*ergo*, there are rights.<sup>99</sup>

Nagel reflects on this "curious type of argument," conceding that its form is not "in general a cogent form of argument."<sup>100</sup> Yet, even if the argument form is invalid outside of moral theory, "It may be suitable to argue that one morality is more likely to be true than another, because the former makes for a better world than the latter--not instrumentally, but intrinsically."<sup>101</sup>

When Nagel writes that one morality might make for "a better world" than another, it typically seems that he means "better" in a collective agent-relative sense, i.e., better *for us*.<sup>102</sup> A

---

<sup>99</sup> Thomas Nagel (1995, 92).

<sup>100</sup> Ibid.

<sup>101</sup> Ibid.

<sup>102</sup> Lippert-Rasmussen reads Nagel as running together two kinds of argument, one involving "the value-for-us" and the other "the value-of-outcome" (135).

morality that includes intrinsic rights is better for us, because it is better to have intrinsic rights than instrumental rights. Intrinsic rights themselves are better, because the kind of inviolability constitutive of intrinsic rights is better than the normatively possible violability constitutive of rights instrumentally construed.

As I mentioned, Nagel does not try to explain any of this. But my framework easily accommodates the value-truth inference regarding personal rights. Rights (and persons) have certain evaluative features, including that they appear extremely valuable to have (and to be). We could be wrong about this; perhaps the best theory of rights, all things considered, involves making some sacrifices with respect to the status of personhood—this would be a regrettable truth. But even if so, if a theory of rights that accommodates inviolability is a theory that makes persons and their rights more valuable, then, to that extent, it is more theoretically plausible than its rivals. Instrumentalist theories of rights, if true, would be true in spite of the theoretical cost they incur by rejecting inviolability. So there is, on reflection, nothing mysterious about the appeal of Nagel's value-truth inference. It is simply indicative of the fact that normative moral theorists are responsible for accommodating not only non-evaluative aspects of morality, but evaluative aspects as well.

Strictly speaking, my account is not a proper rival to Nagel's, because, as I said, Nagel does not really offer an account; he merely offers the apparent plausibility of the inference, leaving the account as basically mysterious.

#### 4.2 *David Enoch*

In his short paper, "Wouldn't it Be Nice if  $p$ , Therefore,  $p$  (for a moral  $p$ )," David Enoch attempts a very different sort of vindication of Nagel's inference. Enoch begins his defense of

Nagel with the observation that, for at least one logical operator, *It is possible that*, the following form of inference is valid:

$$\Diamond\Diamond p \rightarrow \Diamond p$$

If it is possible that it is possible that  $p$ , then it is possible that  $p$ . Enoch observes that for other operators, for example, *I desire that*, the inference rule does not work. I may desire that I desire something and yet not desire it (in other words, second-order desires do not entail first-order desires). Now consider the moral operator, *It is good that*. Suppose we are wondering whether it is good that the world exists, and we could (for some odd reason) directly establish *only* that it would be good if it were good that the world exists. Enoch thinks that it may follow from this that it is, in fact, good that the world exists. Here is how Enoch applies his suggestion to Nagel's article:

If ... the debate over our dignified status can be understood as a debate over something of the form *it is good that  $q$*  (for some  $q$ ), then the inferential move from *it would be good if we had a non-consequentialist dignified status* to *we have such a status* could be vindicated as an instance of  $GGq \rightarrow Gq$ .<sup>103</sup>

But it is not at all clear how the proposition that *we have a non-consequentialist dignified status* is equivalent in the right ways to *it is good that  $q$* , for some  $q$ . Enoch has only shown that, for some operators, if the conclusion of a property-truth inference is itself a proposition falling under the operator, then inferences of the relevant form may go through. But, in the moral case, these are really value-value inferences, not value-truth inferences. So, it is not clear to me that Enoch has even so much as suggested a way of vindicating Nagel's inference, the conclusion of which is not a statement involving a moral-modal operator.

---

<sup>103</sup> David Enoch (200, 223).

#### 4.3 Rob Van Someren Greve

In “Wishful Thinking in Moral Theorizing: Comment on Enoch,” Rob Van Someren Greve criticizes Enoch’s argument for a different reason. Assume both that “it would be good if the actual world were better than it is”, and that a better world “contains more valuable states of affairs than any world that is less good.”<sup>104</sup> Van Someren Greve thinks that this means that “the higher the number of good states of affairs that the actual world contains, the better this world is.”<sup>105</sup> If this is right, then he says that, for any state of affairs that the world contains, it is good if that state of affairs is good. After all, the more good states of affairs there are, the more good the world as a whole is. But if Enoch is right about the operator, *It is good that*, then he is committed to every state of affairs being good, and consequently, he is committed to the world being as good as it can be, given the states of affairs it contains.

I happen to disagree with the contention that, for any state of affairs, it is good if that state of affairs is good. Elsewhere, I defend the idea that it is good, not only that some things are good, but that it is good that some things are bad.<sup>106</sup> A world in which, for example, betrayal is morally indifferent is worse than a world in which betrayal is bad. I take this to be the intuitive force behind an objection to a crude form of Divine Command Theory, according to which moral obligations are whatever God contingently commands. The objection is that, if the crude Divine Command Theory is true, then God could command various horrible things, and such commands would make those things good. Apart from the thought that various horrible things *could not* be good, it seems like it would be *bad* if they were. A crude Divine Command Theory

---

<sup>104</sup> Rob Van Someren Greve (2011, 448).

<sup>105</sup> Ibid.

<sup>106</sup> See Chapter 3, in particular the section on the badness of moral nihilism.

paints a *morally ugly* picture of what moral obligations are like. So, although I think Enoch's argument is a *non-sequitur* with respect to Nagel's point, I nevertheless reject the inferences Van Someren Greve makes from his assumptions about states of affairs to the premises in his *reductio* against Enoch.

#### 4.4 Ryan Preston-Roedder

In his paper, "A Better World," Ryan Preston-Roedder defends what he calls the "better world argument." This argument "rests on the view that morality cannot undermine the aims that it assigns us, together with the view that morality assigns us the aim of bringing about a better world."<sup>107</sup> If those two claims are right, and it turns out that the truth of some moral theory T1 would make for a worse world than the truth of some other moral theory, T2, then it seems that there is a valid value-truth inference available in favor of the truth of T2.

If Preston-Roedder is right, then statements of the form, "Regrettably, *p*," carry with them some tension for any *p* that involves a specific kind of normative moral content. Supposing that it is true that morality, construed realistically, is better than morality, construed antirealistically, then there is at least a sense in which the truth of a metaethical theory might undermine the aims that morality assigns to us. However, this would not obviously mean that a moral theory is subject to the kind of *self-defeat* Preston-Roedder envisages. Preston-Roedder is thinking of moral theories that contain certain substantive principles, like the principle that we should act always to maximize the good, that in themselves make for a worse world and thus defeat one of our moral aims—despite the fact that the very same moral theory (in virtue of being a moral theory) demands that we have this aim. And the fact that a moral theory makes a demand that is

---

<sup>107</sup> Ryan Preston-Roedder (2014, 640).

such that it—just in being a demand—undermines one of our moral aims, is itself a reason to reject that theory.

Metaethical theories like moral realism do not contain substantive principles in this sense, though they may contain accounts of moral phenomena that undermine the aim of making the world better. For example, contrast a realist and a subjectivist reading of what it is for a victim of a wrong to have a claim against a wrongdoer. A realist account of what this means, for example, that the victim is vindicated in a third-personal manner by the external, response-independent system of moral truths, seems to make for a better state of affairs than a simple subjectivist account. If this is true, then adopting a subjectivist position in metaethics might undermine one of our moral aims in a way that moral realism does not. But it would not be a kind of *self*-defeat, since the metaethical theories themselves do not make substantive demands.

Ultimately, I think that Preston-Roedder's account of how a moral theory might be self-defeating with respect to one of our substantive moral aims is less plausible than the considerably less exotic suggestion that moral theorists—both normative and metaethical—are responsible for accommodating, among other things, the evaluative aspects of morality. Taking one of Preston-Roedder's primary examples, it does seem like a crude consequentialist principle that makes no allowances for non-moral value, special obligations to family and friends, meaningful personal projects, and the like, "would deprive us of a desirable form of autonomy that we would have if we were often permitted to promote our private aims."<sup>108</sup> But we can recast this point as one about the *value* of our autonomy, which is something that any moral theorist should strive to accommodate. The better a moral theory accommodates autonomy and its value, the more plausible it is in that respect. If our autonomy just means our freedom to select what the crude

---

<sup>108</sup> Ibid., p. 643.

consequentialist principle dictates, then it is not as valuable as we pre-theoretically take it to be. The consequentialist principle is accordingly less plausible.

The same goes for the kind of independent vindication against wrongdoing I appealed to in contrasting realism with subjectivism. To have the moral verdict in your favor is something that we take to be extremely valuable. If moral verdicts turn out, on a metaethical theory, to be not particularly desirable, then this is a theoretical cost to that theory, albeit not a decisive one.

#### 4.5 *Geoffrey Sayre-McCord*

Finally, in his widely circulated manuscript, “On the Distinction Between a Better Moral Theory and a Theory of a Better Morality,” Geoffrey Sayre-McCord argues that, if we discover that our best theory of morality does not align with our theory of the best morality, then this discovery counts against what we thought was our best theory.<sup>109</sup> His basic idea is that, when value-truth inferences work, they work in virtue of the fact that the “standards of morality” must “justifiably mak[e] a claim on our allegiance.”<sup>110</sup> Ultimately, Sayre-McCord thinks that the point generalizes to all normative concepts. “[C]ompetence with our normative concepts, which purport to establish reasons, commits us to seeing arguments that they do not as a challenge to our understanding of the concepts in question.”<sup>111</sup> In another paper, Sayre-McCord applies these claims to the case of moral dilemmas, arguing that any moral theory that allows for moral

---

<sup>109</sup> Geoffrey Sayre-McCord, “On the Distinction Between a Better Moral Theory and a Theory of a Better Morality,” ms, p. 8.

<sup>110</sup> Ibid., p. 9.

<sup>111</sup> Ibid., p. 13.

dilemmas describes a worse morality (all else equal) than a moral theory that would preclude them, and this is itself a reason to reject the former.<sup>112</sup>

How might this approach apply to the metaethical case (restricting again to metaethical moral theory)? Suppose, with Sayre-McCord, that the standards of morality demand our allegiance. If it is better that those demands be realistic rather than antirealistic, then perhaps moral demands, realistically construed, deserve our allegiance more than moral demands, antirealistically construed. We might recast this as the claim that moral realism better accommodates the authority of morality, where *being authoritative* is understood to be an evaluative feature of morality itself.

Something about the metaethical application of Sayre-McCord's framework seems fishier than the normative application. To see the problem, consider again a metaethical Divine Command Theory on which God's commands are necessarily identical to our moral obligations. Perhaps, moral demands that constitute divine commands deserve our allegiance *even more* than moral demands that do not, in virtue of their being moral demands *plus* commands of God. But this would be a very odd reason to suppose that moral demands *are*, in fact, divine commands. If this were an appropriate inference, then we could infer from the fact that divinely commanded norms are morally better than non-divinely commanded norms, the even more ambitious conclusion that God exists—since the existence of divine commands entails the existence of God.

The only way that such an inference might go through is if we are using the framework I have suggested above, according to which metaethical theorizing must accommodate the evaluative features of morality. If we grant in advance that moral norms are authoritative in a certain way (perhaps, that they demand our allegiance), and it turns out that one metaethical

---

<sup>112</sup> Geoffrey Sayre-McCord (ms), "A Moral Argument Against Moral Dilemmas."

theory best accommodates this fact, then that works in favor of the metaethical theory. But again, this is an ordinary kind of theoretical success—it does not speak to the special nature of normative concepts. To reiterate, none of this means that Sayre-McCord's account is not insightful or illuminating; it is just not necessary for the particular explanatory purpose of vindicating value-truth inferences.

## §5 Conclusion

In this paper, I have addressed an apparently puzzling fact. While inferences of the form, *It would be good that p, therefore p*, are often bad inferences, they seem to have some appeal in some areas of inquiry. Popular objections to the general form of inference are too quick, since such inferences might be licensed in a number of ways. However, once the ways of licensing such inferences are made clear, their successful instances are much less philosophically interesting or surprising than they first appear. But the principled boredom induced by my account is a feature, not a bug. Finally, I surveyed some recent work on these inferences in moral theory, noting how the work relates to and differs from what I say here. The point of the final section was not to say that the comparatively exotic proposals on offer have nothing to contribute to our understanding of morality or normativity, but that they are not necessary for explaining the appeal of value-truth inferences in particular.

## REFERENCES

- Adams, Robert. 1979. "Theories of Actuality." *Nous* 8, no. 3: 211-231.
- Adams, Robert. 1987. "A Modified Divine Command Theory of Ethical Wrongness." In *The Virtue of Faith and Other Essays in Philosophical Theology*, Oxford: Oxford University Press, pp. 97-122.
- Adams, Robert. 1999. *Finite and Infinite Goods: A Framework for Ethics*. Oxford: Oxford University Press.
- Anscombe, Elizabeth. 1958. "Modern Moral Philosophy." *Philosophy* 33 (124): 1-19.
- Blackburn, Simon. 1973. "Moral Realism." In *Morality and Moral Reasoning*, ed. J. Casey. Methuen, London.
- Blackburn, Simon. 1999. "Is Objective Moral Justification Possible on a Quasi-Realist Foundation?" *Inquiry* 42, 2: 213-227.
- Chomsky, Noam. 1988. *Language and Problems of Knowledge: The Managua Lectures*. Cambridge, MA: The MIT Press.
- Cuneo, Terence and Russ Shafer-Landau. 2014. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies* 171: 399-443.
- Dreier, James. 2004. "Meta-Ethics and the Problem of Creeping Minimalism." *Philosophical Perspectives* 18: 23-44.
- Driver, Julia. 2017. "Contingency and Constructivism." In Simon Kirchin (ed.), *Reading Parfit: On What Matters*. New York, NY: Routledge.
- Driver, Julia. 2014. "Gertrude Elizabeth Margaret Anscombe", *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2014/entries/anscombe/>](https://plato.stanford.edu/archives/win2014/entries/anscombe/).
- Dworkin, Ronald. 2013. *Religion Without God*. Cambridge, MA: Harvard University Press.
- Enoch, David. 2009. "Wouldn't It Be Nice If *p*, Therefore *p* (for a moral *p*)," *Utilitas* 21:2, pp. 222-224.
- Erdur, Melis. 2016. "A Moral Argument Against Moral Realism." *Ethical Theory and Moral Practice* 19: 591-602.
- Fodor, Jerry. 1989. "Making Mind Matter More." *Philosophical Topics* 27, no. 1: 59-79.

- Fodor, Jerry. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: The MIT Press.
- Frankfurt, Harry. 2004. *The Reasons of Love*. Princeton University Press.
- Frankfurt, Harry. 1971. "Freedom of the Will and the Concept of a Person," *The Journal of Philosophy*, 68:1, pp. 5-20.
- Hayward, Max. ms. "Immoral Realism."
- Hirsch, Eli. 2009. "Diabolical Mysticism, Death, and Skepticism." *Philosophic Exchange* 39: 2-27.
- Hofweber, Thomas. Ms. "Are there Completely Ineffable Aspects of Reality?"
- Jackson, Frank. 1982. "Epiphenomenal Qualia," *The Philosophical Quarterly* 32 (127): 127-136.
- James, William. 1997. *The Varieties of Religious Experience*. New York, NY: Touchstone.
- Johnston, Mark. 2009. *Saving God*. Princeton, NJ: Princeton University Press.
- Joyce, Richard. 2016a. "Moral Anti-Realism", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/moral-anti-realism/>.
- Joyce, Richard. 2016b. *Essays in Moral Skepticism*. Oxford: Oxford University Press.
- Kahane, Guy. 2016. "If Nothing Matters." *Nous* 50, no. 2: 1-27.
- Kahane, Guy. 2012. "The Value Question in Metaphysics." *Philosophy and Phenomenological Research* 85, no. 1:27-55.
- Kahane, Guy. 2011. "Should We Want God to Exist." *Philosophy and Phenomenological Research* 82, no. 3: 674-696.
- Kamm, F.M. 2001. "Constraining Rights and the Value of Status," in F.M. Kamm (ed.), *Morality, Morality Volume II: Rights, Duties, and Status*. Oxford University Press: 259-289.
- Killoren, David. 2016. "Robust Moral Realism: An Excellent Religion," *International Journal for Philosophy of Religion* 79: 223-237.
- Killoren, David. 2016. "Why Care About Moral Fixed Points?" *Analytic Philosophy* 57 (2): 165-173.

- Kraay, Klaas (ed). 2018. *Does God Matter? Essays on the Axiological Consequences of Theism*. New York, NY: Routledge.
- Leftow, Brian. 2010. "Swinburne on Divine Necessity," *Religious Studies* 46: 41-162.
- Leslie, John. 1979. *Value and Existence* (Basil Blackwell).
- Lippert-Rasmussen, Kasper. 2005. *Deontology, Responsibility, and Equality*. Dissertation. Copenhagen: Department of Media, Cognition and Communication.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Manne, Kate. 2014. "Internalism About Reasons: Sad but True?" *Philosophical Studies* 167: 89-117.
- Mavrodes, George. 1986. "Religion and the Queerness of Morality." In William Wainwright and Robert Audi (eds.), *Rationality, Religious Belief, and Moral Commitment*. Cornell University Press: 213-226.
- McGinn, Colin. 1989. "Can We Solve the Mind-Body Problem?" *Mind* 98, 391: pp. 349-366.
- Mulgan, Tim. 2017. "Beyond Theism and Atheism," *Philosophy Compass* 12(6).
- Nagel, Thomas. 2007. "The Value of Inviolability." In Paul Bloomfield (ed.), *Morality and Self-Interest*. Oxford University Press: 102-114.
- Nagel, Thomas. 1995. "Personal Rights and Public Space," *Philosophy & Public Affairs*, 24:2, pp. 83-107.
- Nietzsche, Friedrich. 2001. *The Gay Science*, ed. Bernard Williams, trans. Josefine Nauckoff. Cambridge: Cambridge University Press.
- Nolan, Daniel. 1997. "Impossible Worlds: A Modest Approach." *Notre Dame Journal for Formal Logic* 38, no. 4: 535-572.
- Olberding, Amy. 1997. "Mourning, Memory, and Identity: A Comparative Study of the Constitution of the Self in Grief." *International Philosophical Quarterly* 37(1): 29-44.
- Olinder, Francén. 2016. "Some Varieties of Metaethical Relativism." *Philosophy Compass* 11 (10): 529-540.
- Olinder, Francén. 2012. "Moral and Metaethics Pluralism: Unity in Variation." *Southern Journal of Philosophy* 50 (4):583-601.
- Parfit, Derek. 2011. *On What Matters* (vol. 2). Oxford: Oxford University Press.

- Parfit, Derek. 2004. "Why Anything? Why This?" in *Metaphysics: A Guide and Anthology*, edited by Tim Crane and Katalin Farkas (Oxford University Press).
- Parker, Theodore. 1853. "Of Justice and Conscience," in *Ten Sermons of Religion* (Boston: Crosby Nichols and Company).
- Paul, L.A. 2016. *Transformative Experience*. Oxford: Oxford University Press.
- Preston-Roedder, Ryan. 2014. "A Better World," *Philosophical Studies* 168, pp. 629-644.
- Rosen, Gideon. ms. "Normative Necessity."
- Sartre, Jean-Paul. *Being and Nothingness*.
- Sayre-McCord, Geoffrey. ms(a). "On the Distinction Between a Better Moral Theory and a Theory of a Better Morality."
- Sayre-McCord, Geoffrey. ms(b). "A Moral Argument Against Moral Dilemmas."
- Street, Sharon. 2017. "Nothing 'Really' Matters, But That's Not What Matters." In Peter Singer (ed.), *Does Anything Really Matter? Essays On Parfit On Objectivity*. Oxford: Oxford University Press.
- Street, Sharon. 2016. "Objectivity and Truth: You'd Better Rethink It." In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* vol. 11.
- Sturgeon, Nicholas. 1986. "What Difference Does it Make Whether Moral Realism is True?" *The Southern Journal of Philosophy* 24: 115-141.
- Swinburne, Richard. 2004. *The Existence of God* (2<sup>nd</sup> ed.). Oxford: Oxford University Press.
- Temkin, Larry. 2017. "Has Parfit's Life Been Wasted? Some Reflections on Part Six of *On What Matters*." In Peter Singer (ed.), *Does Anything Really Matter? Essays On Parfit On Objectivity*. Oxford: Oxford University Press.
- Thucydides. 1996. *History of the Peloponnesian War*, ed. by Robert B. Strassler, trans. by Richard Crawley. New York, NY: Touchstone.
- Van Someren Greve, Rob. 2011. "Wishful Thinking in Moral Theorizing: Comment on Enoch," *Utilitas* 23:4, pp. 447-450.
- Williams, Bernard. 1986. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.

Wolf, Susan. 2010. *Meaning in Life and Why it Matters*. Princeton, NJ: Princeton University press.

Wolf, Susan. 1982. "Moral Saints." *The Journal of Philosophy* 79(8): 419-439.

Wonderly, Monique. 2016. "On Being Attached," *Philosophical Studies* 173: 223-242.

Wright, Crispin. 1988. "Realism, Anti-realism, Irrealism, Quasi-Realism." *Midwest Studies in Philosophy* 12: 25-49.

Wright, Crispin. 1995. "Truth in Ethics." *Ratio* 8, 3: 209-226.

Zagzebski, Linda. 1990. "What if the Impossible Had Been Actual?" In *Christian Theism and the Problems of Philosophy*, ed. by M. Beaty. South Bend, IN: University of Notre Dame Press.