

[Work in progress. Please do not cite or quote without permission.]

Objectivity and Truth: You'd Better Rethink It

Sharon Street
New York University
August 30, 2007

1. Introduction

As human beings, we are capable of viewing ourselves and our values from two very different standpoints. On the one hand is what I'll call, borrowing Kantian language, the *practical standpoint*. When we occupy this standpoint, we understand ourselves as beings who are capable of recognizing what practical reasons we have, what we should or ought to do, what is good, valuable, and worthwhile, what is morally right and wrong, and so on. We think of ourselves, in other words, as beings whose normative judgments (many of them, anyway) are true.¹ For example, when I occupy the practical standpoint, I judge that I have reason to exercise regularly, to contribute to UNICEF, and to balance work and family in certain ways; I think that these and many other judgments about my reasons are true.² On the other hand is what I'll call, again borrowing Kantian language, the *theoretical standpoint*. When we occupy this point of view on ourselves and our values, we understand ourselves as beings who are part of the world of cause and effect, and whose normative judgments are subject to causal explanation. When I occupy the theoretical standpoint, for example, I understand my normative judgments—including my judgments that I have reason to exercise, to contribute to UNICEF, and so on—as having been shaped by

¹ In this paper, I leave aside epistemic reasons and focus exclusively on the case of practical reasons. While I think that the account of practical reasons I defend here ultimately carries over to the case of epistemic reasons, that topic is too much to pursue in this paper. Note also that throughout the paper, I use the word *reason* in the sense of a *normative reason*—in other words, in the sense of a consideration that counts in favor of, or justifies, some action.

² In this paper, I am understanding truth as minimal, so that 'It's true that I have reason to exercise' is just another way of saying 'I have reason to exercise.' Thus, my claim that "we think of ourselves...as beings whose normative judgments...are true" is uncontroversial, even to expressivists such as Simon Blackburn and Allan Gibbard, since on this minimal understanding of truth, anyone who makes any normative judgments at all ipso facto regards some of his normative judgments as true. On minimalism, see Paul Horwich, *Truth*, 2nd ed. (Oxford: Oxford University Press, 1998). On expressivism and minimalism, see Allan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2003).

causes such as my upbringing, cultural background, and inherited psychological tendencies; it is clear to me that had some or all of these factors been different—had I been raised in medieval Japan, for example—I wouldn't have made the same set of normative judgments that I now make.³

Each of these two standpoints on ourselves and our values is inescapable. The practical standpoint is inescapable because we cannot avoid asking and answering questions about our normative reasons. Even if our thinking about such matters is often only implicit and superficial, we cannot avoid considering what to do later on today, or tomorrow morning when we wake up. The theoretical standpoint is equally inescapable, for we cannot avoid seeing ourselves, including our judgments about our reasons, as embedded in the world of cause and effect. We are able to step back from our normative judgments and make them the objects of our attention, just as we're able to do with anything else in the world. When we do so, we cannot avoid seeing these judgments as psychological states whose causal origins and effects we may investigate and discover.

Since both standpoints are inescapable, we must hope that they do not conflict. We must hope that we can coherently view our normative judgments both as true *and* as subject to causal explanation. This is not as easy as one might think, however. In this paper, I will argue that if one is a realist about reasons and value, then these hopes of harmony cannot be met. If one is a realist, then there is, below the surface, a severe and irreconcilable tension between the practical and the theoretical points of view on ourselves—between our understanding of our normative judgments as true and our understanding of them as things with scientifically discoverable causal origins. In contrast, if one is an antirealist about reasons and value, then there is no such conflict. The search for harmony between the two standpoints thus drives us to antirealism.

³ While I borrow Kantian terminology in talking about the practical and theoretical standpoints, I do not borrow any controversial Kantian apparatus. As I employ it, talk of the “practical standpoint” commits us to nothing that we aren't already committed to every time we think we have a reason to do something, such as a reason to put on our socks in the morning. Similarly, talk of the “theoretical standpoint” commits us to nothing that we aren't already committed to every time we think one of our normative judgments (or anything else in the world) has a causal explanation. To occupy the practical point of view is merely to think that one has reasons for action, and to occupy the theoretical point of view is merely to see things (including one's own normative judgments) as subject to causal explanation.

2. Realism and antirealism

As I'll be understanding it, *normative realism* is the view that there are at least some normative facts or truths that hold independently of all our evaluative attitudes, such that an agent can have reason to *Y* even though the conclusion that she has this reason in no way follows from within her own practical point of view, understood roughly as her own set of evaluative attitudes.⁴ *Normative antirealism*, in contrast, is the view that there are *no* normative facts or truths that hold independently of all our evaluative attitudes in this way. According to the antirealist, if an agent has reason to *Y*, then this conclusion must somehow follow from within her own practical point of view; if the conclusion that she has reason to *Y* is *not* entailed from within the standpoint constituted by her own set of evaluative attitudes, then she does not have that reason.

To clarify the distinction, it is helpful to consider the case of an ideally coherent Caligula.⁵ Caligula judges that he has reason to torture others for fun, and this conclusion follows perfectly from within his own practical standpoint; he is utterly consistent in holding that he has this reason, and he is making no mistakes whatsoever about matters of non-normative fact. According to a normative realist, Caligula is nevertheless wrong about his reasons: Caligula does *not* have reason to torture others for fun, even if his entire emotional being delights in such torture, even if he is unimpeachably consistent in his normative opinions, and even if he is in error about no questions of non-normative fact. In other words, the realist holds that there are some reasons—such as the reason not to torture others for fun—that we have quite independently of our evaluative attitudes and what does or doesn't follow from within the standpoint constituted by them.

⁴ *Evaluative attitudes* I am understanding to include states such as desires, attitudes of approval and disapproval, unreflective evaluative tendencies such as the tendency to experience *X* as counting in favor of or demanding *Y*, and consciously or unconsciously held normative judgments about what practical reasons we have, what we should or ought to do, what is good, valuable, and worthwhile, what is morally right and wrong, and so on.

⁵ The example is Allan Gibbard's in "Morality as Consistency in Living: Korsgaard's Kantian Lectures," *Ethics* 110 (October 1999): 140-164, p. 145.

The antirealist, in contrast, denies this. There are two importantly different forms that the antirealist's denial may take. One form of antirealism—what I'll call *Kantian antirealism*—holds that if an agent has reason to *Y*, then the conclusion that he has this reason must follow from within the agent's own practical point of view, but then goes on to make ambitious universalist claims about what may be shown to follow in this way. A Kantian antirealist might hold, for example, that it follows from within every agent's practical point of view that one has reason not to torture others for fun. Such an antirealist in effect denies that an ideally coherent Caligula is possible, arguing that contrary to what's been stipulated about the case, one cannot be perfectly consistent in one's normative judgments, in error about no matters of non-normative fact, *and* judge oneself to have reason to torture others for fun; on the contrary, no matter what one's starting set of evaluative attitudes, it is entailed from within the standpoint constituted by those attitudes that one shouldn't torture others for fun. Christine Korsgaard's view in *The Sources of Normativity*, as I read it, is an example of this brand of antirealism. According to Korsgaard, if one makes any normative judgments at all, one is thereby committed to valuing the humanity of others in such a way that rules out torturing them for fun, among many other forms of treatment.⁶

I call such forms of antirealism *Kantian* because of their bold substantive claims about what follows from within every agent's practical standpoint. According to versions of what I'll call *Humean antirealism*, in contrast, characters like an ideally coherent Caligula are entirely possible, and were such a person ever to exist in real life, he would indeed have reason to torture others for fun. A Humean antirealist agrees with the Kantian that one has reason to *Y* only if this conclusion follows from within the standpoint of one's own set of evaluative attitudes, but then, in opposition to the Kantian, is skeptical of attempts to show that strong, substantive conclusions about our reasons are entailed from within every practical point of view. Bernard Williams's view on the nature of reasons in "Internal and External

⁶ Christine M. Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996).

Reasons” is an example of Humean antirealism, and my own view (with an important qualification that I discuss in section 13) is as well.⁷

We may summarize realism and the two forms of antirealism this way:

	<i>Are there some reasons that agents have even though the conclusion that they have this reason in no way follows from within their own practical point of view?</i>	<i>Is an ideally coherent Caligula possible?</i>	<i>If an ideally coherent Caligula existed, would he have reason to torture others for fun?</i>
<i>Realism</i>	Yes.	Yes.	No, he’d be badly mistaken about his reasons.
<i>Kantian antirealism</i>	No.	No. One is guilty of an incoherence if one judges that one has reason to torture others for fun.	The question has a false presupposition.
<i>Humean antirealism</i>	No.	Yes.	Yes.

The central focus of this paper is the realism/antirealism debate, so understood. The intramural dispute between Kantian and Humean versions of antirealism will be relevant only briefly at a couple of points, and I will make no attempt to settle it here.⁸

3. *The paper’s strategy*

In his 1996 article “Objectivity and Truth: You’d Better Believe It,” Ronald Dworkin defends a version of normative realism that in my view has not yet received an adequate response from antirealists.⁹ One of the main goals of the present paper is to offer such a response. As will become clear, there is a great deal about which I think Dworkin is exactly right. Indeed, my strategy in the paper is to adopt wholesale almost all of the major points that Dworkin argues for in “Objectivity and Truth: You’d Better Believe It”—but with one major exception. The exception is this: I will argue that Dworkin’s

⁷ Bernard Williams, “Internal and External Reasons,” in *Moral Luck* (Cambridge: Cambridge University Press, 1981), 101-113. I explain my own view in more depth below. For a fuller statement of my positive view on the nature of reasons, see “Constructivism About Reasons,” forthcoming in *Oxford Studies in Metaethics*, vol. 3, ed. Russ Shafer-Landau (Oxford: Clarendon Press, 2008).

⁸ I discuss this intramural dispute at more length in “Constructivism and the Foundations of Our Reasons” (draft).

⁹ Ronald Dworkin, “Objectivity and Truth: You’d Better Believe It,” *Philosophy & Public Affairs* 25 (1996): 87-139. All page references in the text are to this work unless otherwise noted.

endorsement of *realism*—understood exactly as he himself wishes to understand it, namely as an “internal” normative claim—involves a stark refusal to integrate the practical point of view (from which we make judgments about what we have reason to do), with the theoretical point of view (from which we inquire into the causal origins of, among other things, our judgments about what we have reason to do), and that this refusal renders the position deeply unstable. Working from normative premises that Dworkin himself accepts—and which indeed I think anyone who engages in practical reasoning at all must accept—I will try to show why we are all forced to antirealism. We are all forced to antirealism, that is, I’ll argue, unless we are content to rest our realism on nothing more than an unreasoned faith, with realism about reasons and value thus becoming a rather odd form of religion.

While responding to Dworkin’s position is important in itself, my case study of Dworkin’s view has three larger goals. The first is to offer a skeptical antirealist argument that applies not only against Dworkin’s view, but also against all versions of an increasingly popular brand of normative realism—what is often called *non-naturalist realism*, and whose contemporary defenders include, in addition to Dworkin himself, David Enoch, William FitzPatrick, Thomas Nagel, Derek Parfit, T. M. Scanlon, and Russ Shafer-Landau.¹⁰ As realists, these philosophers all endorse the basic realist conviction that there are some reasons—for instance, the reason not to torture others for fun—that we have quite independently of our evaluative attitudes and what does or doesn’t follow from within the standpoint constituted by them. In addition to this basic tenet of realism, non-naturalist realists also hold that it is a mistake to understand normative facts or truths as natural facts or truths of the kind investigated by the natural and social sciences; these realists think that normative facts are not properly understood as things with causal powers, and that it is a mistake to think that the test for whether normative facts exist is whether we need to suppose them in our best causal explanations of anything. All of these latter claims they hold in opposition to *naturalist* versions of normative realism, whose proponents likewise accept the basic tenet of realism, but then argue that normative facts *are* best understood as natural facts, that normative facts *do*

¹⁰ [References.]

possess causal powers, and that it is *correct* to think that the test for whether normative facts exist is whether we need to suppose them in our best causal explanations of anything. Examples of such naturalist realists include David Brink and Nicholas Sturgeon.¹¹

The second larger goal of the paper is to respond to some important objections to normative antirealism. As I have mentioned, my own view is antirealist, and I think the resources and plausibility of antirealism are often seriously underestimated. The two most important objections I address are the objection that antirealism's consequences are so repugnant and deeply counterintuitive that antirealism loses out in a holistic comparison with realism (section 12), and the objection that antirealism is ultimately self-defeating because of the way in which it applies to itself (section 13).

The third and final larger goal of the paper is to sketch, in broad brushstroke, a completely general argument against normative realism, whether of a non-naturalist or naturalist variety (or even of a quasi-realist variety, as I discuss elsewhere¹²). According to this general argument, realism about normativity—not just Dworkin's version, or non-naturalist versions, but any version—involves an unacceptable failure to integrate the practical and theoretical standpoints on ourselves—a failure that forces us to antirealism, whether Kantian or Humean. It must be emphasized, however, that this sketch of a general argument is just that—a mere sketch, for there is not space here to defend at least two of its key premises—namely that naturalist versions of normative realism fail, and that antirealism in some form or other can ward off all major objections well enough to defeat realism when we step back and compare realism and antirealism from a broad, holistic perspective.¹³

¹¹ David S. Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), and Nicholas Sturgeon, "Moral Explanations," in *Morality, Reason and Truth*, eds. David Copp and David Zimmerman (Totowa, NJ: Rowman and Allanheld, 1985).

¹² I discuss quasi-realism in "Why You Can't Have Your Cake and Eat It Too in Metaethics: A Naturalist Challenge to Quasi-Realism About Normativity" (draft).

¹³ Both of these premises I defend elsewhere. I argue that naturalist versions of normative realism fail in "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies* 127, no. 1 (January 2006): 109-166, especially section 7, and in "Reply to Copp," forthcoming in *Philosophical Issues* 18 (2008). I address some of the most important objections to my own brand of antirealism in "Constructivism About Reasons."

The paper's argument proceeds as follows. After briefly summarizing Dworkin's position in section 4, I identify, in section 5, a puzzle that we all face. In sections 6 and 7, I discuss two possible solutions to this puzzle, both of which Dworkin rejects—the first wrongly, the second rightly, in my view. In section 8, I present Dworkin's own preferred solution to the puzzle, and in section 9 I argue that Dworkin's solution is unacceptable. I gradually develop and refine my argument by responding to four objections in sections 10 through 13, and in section 14, I conclude with an observation about where all of this leaves us.

4. *A brief summary of Dworkin's position*

In "Objectivity and Truth: You'd Better Believe It," Dworkin distinguishes between two kinds of skepticism about the normative domain, *external* and *internal*. The *external skeptic* seeks to establish her skeptical position from a standpoint *outside* the normative domain, in the sense that she attempts to argue, without relying on any normative premises whatsoever, that there is no such thing as normative truth. The external skeptic might begin, for example, with purely non-normative premises about what kinds of properties exist in the world, or with merely anthropological observations about the range and depth of disagreement about normative matters, and from there try to argue that there is no such thing as normative truth (p. 88).

The *internal skeptic*, in contrast, does not attempt to stand completely outside the normative domain. Instead, he works from within it, denying some familiar group of normative claims, but justifying that denial by appeal to some further set of normative claims—perhaps more general or theoretical or counterfactual ones (p. 90). Some versions of internal skepticism are limited, while others are broad. As an illustration of the former, Dworkin points to skepticism about conventional sexual morality. Working from the normative assumption that suffering is the only thing that is inherently bad, for example, the internal skeptic about conventional sexual morality might argue that since many acts condemned by conventional sexual morality lead to no suffering, conventional sexual morality ought to be rejected (pp. 90-91). As an example of a broad version of internal skepticism, Dworkin points to

internal skepticism about morality as a whole. Working from the normative assumption that God is the only plausible basis for moral requirements, for instance, the internal skeptic about morality as a whole might argue that there is no God, and therefore also no moral requirements (p. 91).

The main goal of Dworkin's paper is to argue that external skepticism about the normative domain is always confused. According to Dworkin, "Any successful—really, any intelligible—argument that evaluative propositions are neither true nor false must be internal to the evaluative domain rather than [external] to it" (p. 89). He summarizes:

[N]on-moral discoveries cannot undermine or structurally change morality without morality's help. They furnish information that engages with deep substantive moral judgments, and their impact depends on what those deep judgments themselves declare or assume.... Morality is a distinct, independent dimension of our experience, and it exercises its own sovereignty. We cannot argue ourselves free of it except by its own leave. (pp. 127-128)¹⁴

One way of reading Dworkin's paper is as offering a fancy restatement of what is often called *Hume's Law*, according to which there must always be at least one normative claim among one's premises in order for one validly to derive a normative claim as a conclusion.¹⁵ What makes Dworkin's restatement of Hume's Law illuminating and important is Dworkin's argument that the set of claims that are best understood as normative is much larger than many have thought, encompassing all kinds of claims that one might have thought were "purely metaethical" and neutral with respect to implications about how we should live our lives.¹⁶

For example, when we think of the set of normative claims, the sorts of claims that often leap to mind are claims such as 'Abortion is wrong' or 'I should exercise regularly.' But according to Dworkin, the set of normative claims equally well encompasses more dressed-up claims such as the following:

"What I said [just now about abortion's being wrong] was not just venting my emotions or describing or

¹⁴ When Dworkin states his conclusions he sometimes focuses, as in this passage, on the domain of *morality*. But he is best read as making claims about the domain of the *normative* in general, and this is how I read him throughout this paper (though recall that I am focusing on the domain of practical reasons, setting aside the domain of epistemic reasons—see note 1 above).

¹⁵ David Hume, *A Treatise of Human Nature* (1739-1740), second edition, eds. L. A. Selby-Bigge and P. H. Nidditch (Oxford: Oxford University Press, 1978), Book III, Part ii, Section 1.

¹⁶ Blackburn and Gibbard have made similar arguments [references].

expressing my own or anyone else's attitudes. My opinions are...really and objectively true.... They are universal and they are absolute.... They are reports...of how things really are out there in an independent, subsisting realm of moral facts" (p. 97). With regard to these claims, others like them, and also denials of such claims, Dworkin argues two things: first, that there is a plausible interpretation of each of them that shows each to be a substantive normative claim—a claim that takes a substantive position (however abstract or counterfactual in nature) on how we should live; and second, that we cannot find any plausible interpretation of such claims that shows them to be anything *other* than substantive normative claims.

Dworkin's argument for these two points occupies much of his paper and is too much to summarize in full here, but for our purposes we need consider only one type of example. Consider, in particular, claims about the mind-independence of morality or normativity more generally. All such claims, Dworkin argues, are best understood as substantive normative claims about what reasons we would have under counterfactual conditions in which our attitudes were different.¹⁷ For example, the claim that 'Abortion is objectively wrong' is equivalent to the claim that abortion would still be wrong even under counterfactual conditions in which no one thought it was. And the claim that 'Moral facts are out there in an independent realm'—while this is not something regular people ever say, Dworkin notes—is best understood as an inflated, metaphorical way of insisting that what is right and wrong doesn't vary with people's views and attitudes regarding what's right and wrong.

The important point for our purposes, then, is that Dworkin understands the realism/antirealism debate as I characterized it in section 2 as a substantive normative debate. When the realist and antirealist argue with one another, they are simply disputing what agents' reasons are under various conditions. The only thing that makes their disagreement at all different from any regular normative debate—such as a debate about whether abortion is wrong—is that the realist and antirealist are giving special attention to counterfactual cases in which the agents whose reasons are in question have very different evaluative

¹⁷ This is a point that Blackburn and Gibbard have defended as well.

attitudes than we do (delighting in torture, for instance), and seeking to arrive at exceptionally broad generalizations about how agents' reasons vary or not with their evaluative attitudes. At the end of the day, though, these are mere differences in degree, not kind, on Dworkin's view. Fundamentally, the realist/antirealist debate is no different than any other ordinary normative debate about who should do what under what circumstances.

In this paper, I grant Dworkin all of these points. I assume that realism and antirealism are best understood as substantive normative claims, and I also accept Dworkin's point that the only viable form of normative skepticism is internal—that is, that one can legitimately reach skeptical conclusions only about subparts of the normative domain, and only by working from further normative assumptions. As Dworkin writes, "We can't be skeptical, even about values, all the way down" (p. 89). What I intend to defend, then, is not a thoroughgoing, external skepticism about normativity that relies on no normative premises at all, but rather a more limited, internal skepticism directed specifically at normative realism, understood as a substantive normative position.

5. A puzzle we all face

There is a puzzle that each of us faces in virtue of our capacity to occupy both the practical and the theoretical points of view on ourselves and our values. On the one hand, as I have said, we think that many of our normative judgments are true. While we of course do not view ourselves as infallible about normative matters, we are nevertheless convinced that we are not completely hopeless when it comes to figuring out what we have reason to do. On the other hand, we recognize that causal forces have shaped our normative judgments in countless ways: we recognize that had we experienced a different upbringing, or had we been born into a different culture, or had we lived in a different millennium, or had our species been the product of a different evolutionary history, and so on, we would have made very different normative judgments than the ones we now make.

This raises a puzzle. Insofar as we regard our normative judgments as true, we must agree that there is a striking coincidence between (1) the normative judgments that are true, and (2) the normative

judgments that causal forces led us to believe. This coincidence between the normative truth and what causes led us to believe is puzzling and demands some kind of explanation. Call this puzzle the *practical/theoretical puzzle*.¹⁸

I take it to be obvious that there is *some* puzzle here. Note, however, two things. First, in stating the practical/theoretical puzzle, I do not mean with the use of the word *coincidence* to be presupposing anything beyond the point that the normative judgments that are true and the normative judgments that causal forces led us to believe *coincide* to a large extent. I am not, for example, making any assumptions about whether this coincidence was unlikely, or will be difficult to account for; for all I've said so far, the coincidence could be as mundane and easy to account for as the coincidence between the presence of McDonald's franchises and golden arches on poles, for example—a "puzzle" barely worthy of the name, since it's so easy to explain. Second, although I claim that *some* kind of explanation of the coincidence is demanded, I am presupposing nothing about what *kind* of explanation is needed—whether it be causal, normative, philosophical, or of any other particular kind. Nor am I ruling out the possibility that the explanation will just consist in an account of why no (further) explanation of the coincidence is needed—just as we think no explanation is needed when a day of sunshine coincides with one's birthday.

Dworkin faces the practical/theoretical puzzle just like the rest of us. Just like the rest of us, he takes himself to be capable of recognizing normative truths—ones that hold independently of our

¹⁸ My talk of what causal forces "led us to believe" about normative matters might prompt the objection that I am overlooking or underestimating the role of reflection in shaping our normative judgments. After all, we are not just passively "led to" our normative judgments; very often they emerge from active reflection. I have no quarrel with this point; I agree that active reflection has played a major role in shaping our normative judgments. The key point for my purposes is just the commonsensical one that no matter what my current set of normative views, and no matter how much they are the product of my own and others' reflection, it is nevertheless clear that these views were also strongly shaped by contingencies of upbringing, history, culture, education, genetic inheritance, and so on, such that had these causal factors been very different—endowing me with a different starting set of evaluative attitudes, and acting on me along the way in different ways—then my current normative views would also have been very different. Dworkin fully agrees with this counterfactual point (p. 125), and it is in this sense only that I mean that causal forces have "led" us to our current normative views: I mean that had the causal factors that gave us our starting set of values and influenced our normative views along the way been different, our current normative views would also have been different. For further relevant discussion, see "A Darwinian Dilemma for Realist Theories of Value," section 5.

attitudes, since he is a realist—and he also thinks that there are good causal explanations of why he makes the normative judgments he does (p. 104). He recognizes that had he had a different upbringing, or been born into a different culture, and so on, he would have made very different normative judgments (p. 125). So for him too, there is a puzzling coincidence that needs to be explained.

In the next three sections, I discuss three possible solutions to the puzzle: (1) normative antirealism; (2) normative realism coupled with the claim that we causally interact with the independent normative truth; and (3) Dworkin's preferred solution. Since space is limited, I survey the first two solutions only briefly; in this paper, I wish to focus on Dworkin's proposed solution and its viability.

6. *One possible solution: antirealism*

My own view is that the solution to the practical/theoretical puzzle is normative. In particular, the best explanation of the coincidence in question is antirealism itself, which recall we are understanding as a substantive normative view.

To see how antirealism solves the puzzle, consider constructivism, a specific brand of antirealism. According to the version of constructivism I favor, the fact that X is a reason for agent A to Y is constituted by the fact that the judgment that X is a reason (for A) to Y withstands scrutiny from the standpoint of A 's other judgments about reasons. So, for example, *that it would be fun* is a reason for A to kick dogs if and only if that normative judgment withstands scrutiny from the standpoint of A 's other judgments about reasons; indeed, A 's having the reason is constituted by that fact. This view obviously raises many questions. For example, the notion of “withstanding scrutiny” requires clarification and must receive a purely formal characterization (as opposed to a substantive normative one) if the account is not to be unhelpfully circular. I address this and other objections elsewhere.¹⁹ Two further objections—according to which the view has extremely counterintuitive consequences, and according to which it is self-defeating because it applies to itself—I speak to later in this paper.

¹⁹ In “Constructivism About Reasons.”

For the moment, though, the important point is this. The “puzzling coincidence” we are concerned with is no longer puzzling on this understanding of the nature of reasons and value. If constructivism is true, then *of course* there’s a striking coincidence between true normative judgments, on the one hand, and the normative judgments that causal forces led us to make, on the other, because normative truth just *is* a (rather complex) function of the normative judgments that causal forces led us to make. Causal forces originally led us to affirm normative judgments *X*, *Y*, and *Z*, say, and the true normative judgments are those which withstand scrutiny from the standpoint of judgments *X*, *Y*, and *Z*. There is therefore no mystery about why there is significant overlap between the true normative judgments and the ones that causal forces led us to make, for we can expect the true normative judgments to consist, at least in part, in a pruned and coherent subset of the ones bequeathed to us by causal forces.²⁰

As we have seen, however, Dworkin, in the company of many others, rejects antirealism and embraces realism. As I discuss in more depth in section 12, Dworkin thinks that all things considered, it is unacceptable to hold, as the constructivist does, that our normative reasons are ultimately always a function of what we take them to be. So if, along with Dworkin and many others, we assume normative realism—if we assume that there are normative truths that hold independently of all our evaluative attitudes and what follows from within the standpoint constituted by them—then what other solutions to the practical/theoretical puzzle are available to us? There are two other main options.

7. *A second possible solution: causal interaction with the independent normative truth*

Turn now to a second way of explaining the coincidence between true normative judgments and the normative judgments that causal forces led us to make. One might seek to account for the coincidence with some kind of *causal* explanation. To see the general form of the solution being envisioned, notice how we might raise an analogous puzzle about the case of, say, judgments about our manifest surroundings.²¹ (Such judgments would include, for example, my judgment that there is a tree in front of

²⁰ For further discussion, see “A Darwinian Dilemma for Realist Theories of Reasons,” section 10.

²¹ I borrow talk of our “manifest surroundings” from Allan Gibbard, *Wise Choices, Apt Feelings* (Cambridge, MA: Harvard University Press, 1990), chapter 6.

me, a boulder behind me, and a lake to my left.) After all, we think that our judgments about such matters (many of them, anyway) are true. And yet we also think that these judgments are subject to causal explanation. And so we have all that is needed to generate the puzzle: There is a striking coincidence between the judgments about our manifest surroundings that we think are *true*, and the judgments that causal influences led us to make. This coincidence demands an explanation just as much as the coincidence we have noted in the practical case.

It would be foolish to suppose that an easy explanation is available, and offering a fully satisfactory explanation is not the topic of this paper. Yet it seems to me that we have a pretty good idea of how the explanation goes. It's something along the following lines. The objects of judgments about our manifest surroundings are things with causal powers—things such as trees, boulders, and lakes. Because of their causal powers, these are the sorts of things it promotes reproductive success not to miss: after all, one can run into a tree, be crushed by a boulder, or drown in a lake. Evolution by natural selection—a causal process—thus explains how we came to be able to make true judgments about the presence and absence of these things. And so we see—in very rough outline, anyway—why there tends to be a fair coincidence between the judgments about our manifest surroundings that are true and the ones that causal forces have led us to make.²²

Encouraged by this solution to the puzzle in the case of judgments about our manifest surroundings, one might hope that a similar solution can be offered to the parallel puzzle concerning normative judgments. One example of this solution might be to affirm the existence of what Dworkin calls *morons* (p. 104)—special moral particles with causal powers—and then to explain why an ability to detect morons tended to promote the reproductive success of ancestors who possessed this ability. The thought is that perhaps an inability to detect these particles led to decreased reproductive success, just as an inability to detect boulders, trees, and lakes did.

²² [References and discussion of some complications.]

As his *moron* terminology suggests, however, Dworkin utterly rejects any solution along these lines. He thinks the idea of morons is absurd (pp. 104-105), and indeed takes the view that normative properties *never* play a role in our best causal explanations. Moral rightness and wrongness, goodness, normative reasons, value, and so on, in Dworkin's view, are not things with causal powers at all, and it is a misguided test for the existence of these things to ask whether or not they play a role in our best causal explanations (p. 119). As Thomas Nagel has put the point, "Mackie [has argued that] reasons play no role in causal explanations. But it begs the question to assume that this sort of explanatory necessity is the test of reality for values. The claim that certain reasons exist is a normative claim, not a claim about the best causal explanation of anything."²³ As we saw earlier, this rejection of the idea that reasons and values are things with causal powers is a distinguishing feature of non-naturalist versions of normative realism in general.

I think there is more to be said about this second possible solution to the puzzle than Dworkin's talk of morons suggests. Yet in the end, I agree with Dworkin, Nagel, and other non-naturalists on this point. Rightness, wrongness, goodness, normative reasons, value, and so on are very different things than trees and boulders, and a causal story will not solve the practical/theoretical puzzle. I will not argue this point here, however; I say more about it elsewhere.²⁴ I mention this second possible solution only to set it aside; in what follows I will assume with Dworkin and other non-naturalist realists that this second solution to the puzzle fails, allowing us to focus on Dworkin's proposed solution.

8. *A third possible solution: Dworkin's preferred solution*

At its most basic level, Dworkin's answer to the practical/theoretical puzzle is that the coincidence in question is just that—a coincidence; one ought to understand it as a "piece of luck" (p. 125) that the true normative judgments, on the one hand, and the normative judgments that causes led one to make, on the other, line up as much as they do. Dworkin's view on this topic emerges most clearly in

²³ Thomas Nagel, *The View From Nowhere* (New York: Oxford University Press, 1986), p. 144.

²⁴ "A Darwinian Dilemma for Realist Theories of Value," especially sections 6, 7, and 8.

his discussion of an example. Imagine, Dworkin says, that you (who are among the well-to-do) believe that justice requires higher taxes for redistribution to the poor. Then one day you discover the truth of the “bleak thesis” that all human moral convictions are ultimately motivated by self-interest, and become convinced that you wouldn’t have arrived at your moral view in support of higher taxes had it not somehow (you’re not sure how) promoted your self-interest to do so. About this case, Dworkin writes:

[W]hy shouldn’t you count it as a piece of luck—a special example of what Bernard Williams has called moral luck—that your self-interest and justice here coincide? You realize that you would have had a very different view if your own self-interest had favored that different view instead (just as you realize now, independently of the bleak thesis, that you would have had different views if you had been born into a very different culture). But that fact alone cannot undercut your conviction about justice, and therefore about your moral luck. You have to think something—either that you have an obligation to vote for higher taxes or that you do not—and you have no reason, just in what the bleak thesis declares, to turn your back on all the reasons...that inclined you to think that you do. You lack, so far, in other words, something that neither the bleak thesis nor any other piece of psychology can provide on its own. You lack a normative connection between the bleak psychology and any conclusion about justice, or any other conclusion about how you should vote or act. (p. 125)

In this passage we again see Hume’s Law at work: findings about the causal origins of one’s normative judgments cannot by themselves, without the assistance of further normative premises, have any upshot regarding the truth of those normative judgments. Dworkin grants that in some cases there *will* be a “normative connection” that enables us to move from findings about the causal origins of a given normative view to the conclusion that the view is false (pp. 125-126); in such cases, we’ll have arrived at an instance of internal skepticism about some subpart of the normative domain. Dworkin’s point, though, is that in the *absence* of any such normative connection, we must regard any coincidence between the true normative judgments, on the one hand, and the judgments that causes led us to make, on the other, as mere coincidence, and nothing more.

This view—according to which the coincidence noted by the practical/theoretical puzzle is *mere* coincidence—might strike some as implausible on its face unless the following point is emphasized. While some coincidences (for instance, the coincidence between McDonald’s franchises and golden arches on poles) cannot plausibly be dismissed as mere coincidence, and nothing more, there are plenty of other coincidences in life that we do think are mere coincidences—coincidences regarding which no

further explanation is required, and indeed regarding which we think it would be positively confused to insist upon one. Dworkin clearly regards the coincidence highlighted by the practical/theoretical puzzle as belonging to this latter category—the category of coincidences for which it is positively misguided to think further explanation is necessary. To fill out Dworkin’s view, it is helpful to consider an example.

Suppose that, against all odds, I win the New York State Lottery. Then I am faced with a striking coincidence: the very lottery ticket I bought is the very lottery ticket that was randomly chosen as the winning ticket. In such a case, the right view seems to be this. There are, in the end, just two explanations here: the explanation of how it happened that I bought this particular lottery ticket, and the explanation of how it happened that this particular lottery ticket was the one chosen at random. Assuming there is no evidence that the lottery was anything but fairly administered, then once these two things are fully explained, there is nothing further left to explain; the ticket I bought is the one that was selected, and that’s the end of it. The right thing to say, in the end, is that I was just lucky. If I understand both of the explanations in question, and believe, with no evidence to the contrary, that the lottery was fairly administered, and yet nevertheless I continue to ask, “Yes, but why did *I* win? Why was *I* the lucky one?”, then my question is a confused request for a deeper explanation that doesn’t exist.

Dworkin’s view, as I read it, is that the coincidence in the normative case is similar. There are, in the end, just two explanations here: the *normative* explanation of why normative judgments *X*, *Y*, and *Z* are the true ones, on the one hand, and the *causal* explanation of why causal forces led me to affirm those very same judgments. Assuming there is no reason to think there is any “normative connection” between the causal explanation and the normative conclusions in question, then once these two things are fully explained, there is nothing further left to explain; the normative judgments that are true are the ones the causes led me to affirm, and that’s the end of it. The right thing to say, in the end, is that I was just lucky. If I understand both of the explanations in question, and believe, with no evidence to the contrary, that there is no normative connection between the causal explanation and the normative conclusions in question, and yet nevertheless I continue to ask, “Yes, but how did it happen that the causes led me to the

true normative judgments? Why did I get lucky on this?”, then my question is a confused request for a deeper explanation that doesn’t exist.

Or so I think Dworkin would argue in reply to the practical/theoretical puzzle: the alleged “puzzle” is no more a genuine puzzle than the “puzzle” associated with, say, winning a fair lottery, or getting struck by lightning, or being born.

9. *Why Dworkin’s solution is unacceptable*

There is no question that life is full of coincidences regarding which it is misguided to insist upon a further explanation. But our ability to grasp the normative truth is not one of these. On the contrary, there is a critical disanalogy between the case of winning the New York State Lottery (or getting struck by lightning or being born) and the case of hitting upon the normative truth.

In the lottery case, we are supposing that I already know that I won the lottery—presumably because I got the official call, the press showed up, and the checks have started rolling in. In such a case, the account offered above seems perfectly adequate: after a certain point, you just have the two explanations (the explanation of why this was the ticket I bought and the explanation of why this was the ticket that was randomly chosen), and beyond that, there’s nothing to say except that it was pure luck that I won. The normative case is not like that, however. Whereas in the lottery case, I have plenty of non-trivially-question-begging evidence that I won the lottery (in the form of the phone call, the press, the checks, etc.), in the normative case, by Dworkin’s own admission, we ultimately have no non-trivially-question-begging evidence that our own normative judgments are true. And this makes for a crucial difference.

To see the point here, it is important to move slowly. For of course there is a sense in which we *do* have non-trivially-question-begging evidence that our normative judgments are true. To return to Dworkin’s example, suppose I make the normative judgment that justice requires higher taxes for redistribution to the poor. There are plenty of non-trivially-question-begging normative reasons I can give in favor of this view: perhaps I’ll cite considerations regarding the equality of persons, the

arbitrariness of the class circumstances one is born into, various arguments from Rawls, and so forth. Just as Dworkin argues, merely noticing that I wouldn't have made this judgment about justice had I experienced a different upbringing or been born into a different culture does nothing, by itself, to undermine the force of those reasons. In other words, if we raise the practical/theoretical puzzle about any *individual* normative judgment, then Dworkin's answer is satisfactory as far as it goes.

But the practical/theoretical puzzle as I have presented it is not about an individual normative judgment. Instead, it concerns our normative judgments in general: it notes that whatever our more specific normative views, we are forced to think that there is a *general* coincidence between the true normative judgments and the ones that causal forces led us to make, and demands that this coincidence be explained. It's in response to this demand that Dworkin's proposed solution breaks down.

More specifically: Dworkin grants that there are other possible agents such that they are (in Dworkin's view) dead wrong about normative matters, but such that they are suffering from no cognitive impairment that is manifested in ways beyond their inability to reach the right substantive answers on normative questions (pp. 106-108 and 121-122). Return to the case of the ideally coherent Caligula. Dworkin agrees that when it comes to cases such as this, there will sometimes be no way of saying what mistake the other agent is making except to say that he is wrong about his reasons. As Dworkin writes,

We may be forced to concede, in some cases, that those who held different views lacked no information we have, and were subject to no different distorting influences. All that we can say, by way of explanation of the difference, is that they did not "see" or show sufficient "sensitivity" to what we "see" or "sense," and these metaphors may have nothing behind them but the bare and unsubstantiated conviction that our capacity for moral judgment functions better than theirs did. (pp. 121-122)

Dworkin thus grants that there are other possible agents such that one has no non-trivially-question-begging reason to think that oneself is right while they are wrong. And there seems to be no reason in principle why there couldn't be many—indeed infinitely many—such agents (both real and possible).

But this raises a problem. If one accepts Dworkin's view that the coincidence between the true normative judgments and the judgments that causal forces led one to embrace is best understood as *mere* coincidence, then one must view oneself as the lucky one among all of these agents—the lucky one whose

causal history just happened to be such as to put one in a position to recognize the independent normative truth that all these other agents, due to their unlucky set of starting points, are unable to see. And given that there are innumerable such agents (both real and possible), the odds of having won this “normative lottery” are even lower (to put it mildly) than those of having won the New York State Lottery.

Now of course the mere fact that the odds of something’s happening were phenomenally low doesn’t, by itself, entail that that thing didn’t happen. If we already know that I won the New York State Lottery, or was struck by lightning, or was born, it is obviously silly to argue that these things didn’t happen on the grounds that their occurrence was phenomenally unlikely. But when the odds of something’s happening are phenomenally low *and* one lacks any non-trivially-question-begging evidence that the thing happened, then you should believe that in all likelihood, it didn’t happen. For example, if you have no non-trivially-question-begging evidence that you won the New York State Lottery—no evidence whatsoever in the form of a notification call, or visits from the press, or checks rolling in—then it’s crazy to believe that you did win the lottery; imagine someone who insists that she won, and yet when we ask her how she knows, she merely tells us again, in so many words, “Because I won!” This, however, is exactly one’s situation in the normative case, if we accept Dworkin’s view. By Dworkin’s own admission, you cannot point to any non-trivially-question-begging reason to think that you were the lucky one among the innumerable possible agents whose normative views exhibit no cognitive impairment that is manifested in ways beyond their inability to reach (what you regard as) the right substantive answers on normative questions. But then just as in the New York State Lottery case, it’s crazy to believe that you won this normative lottery. The odds of your having won are phenomenally low *and* you lack any non-trivially-question-begging evidence that you did. So you should believe that in all likelihood, you didn’t—that in all likelihood, you are one of the unlucky ones whose system of normative judgments hopelessly diverges from the normative truth.

It is no surprise, then, that in his discussion of the example in which one learns that self-interest drove one to one’s view about the justice of higher taxes, Dworkin ends up talking about viewing oneself

as among the “elect.” According to Dworkin, if you conclude that, in your case, self-interest and morality indeed coincided, then you will

be in much the same state of mind as the theologians of predestination, from Augustine through Calvin, who supposed that people saw the truth only because they had been elected by God. Those theologians continued to think that they were among the elect, though they knew that people could think that who were not. (p. 126)

My complaint, in short: Dworkin and any other normative realist who accepts his response to the practical/theoretical puzzle must regard themselves as among the “normative elect,” while simultaneously admitting that they have no non-trivially-question-begging reason to think this, and indeed positive reason to think the contrary (namely, the sheer odds against it). But this is just an abandonment of reasoned philosophical argument. Rather than any *answer* to the practical/theoretical puzzle about the coincidence between the normative truth and what causal forces led us to believe, all we have here is an article of faith. Normative realism has become a strange form of religion—a religion stripped clean of everything except the bare conviction that there are independent normative truths that one is capable of recognizing.

To summarize the larger dialectical situation: We have considered three possible solutions to the practical/theoretical puzzle. One solution is naturalist normative realism, which couples normative realism with the view that we causally interact with the independent normative truth. This solution we have set aside for the sake of argument, since Dworkin thinks, and I agree, that the view is unacceptable. I have now argued that Dworkin’s non-naturalist realist solution is also unacceptable. It forces us to choose between the view that we are in all likelihood hopeless as normative judges and an unreasoned faith that we are among the normative elect. The view that we are in all likelihood hopeless as normative judges is (as I’ll argue in more depth in section 12) completely unacceptable. Faith is an option, but the theorist who opts for it in the face of strong positive reason to think that she is hopeless has removed herself from the realm of reasoned philosophical argument. The solution that’s left is normative antirealism. This, in my view, is the solution that the failure of the other options forces us to. While I cannot make a complete case for this larger conclusion here—having left key premises (such as the failure

of naturalist normative realism) undefended—in what follows I’ll continue to develop a key strand of the larger case by responding to four objections to my argument against Dworkin’s solution.

10. First objection: The argument is a misguided request for an “external” justification

To begin with, one might object that my argument is asking for the very thing that Dworkin has already argued is impossible—namely a justification of one’s normative convictions from some standpoint *outside* the normative domain. After all, the argument notes the countless number of internally consistent normative viewpoints, and then might seem to step back from all of them, asking for a *reason* to think that one’s own internally consistent view, as opposed to any of the others, is the right one.²⁵ It might seem that in making this request, I am demanding a justification of one’s normative views that does not itself presuppose any further normative view. Yet this would just be a demand that we derive an “ought” from “is” premises alone—a violation of Hume’s Law, and an attempt to establish the kind of external skepticism that Dworkin has argued—and I have agreed—is misguided.

But this is *not* what I am asking for. The challenge I have issued to Dworkin’s position is thoroughly internal to the normative domain, and it is assumed that the answer, if there is one, will be internal as well. The challenge, after all, is not directed at the normative domain in general; it is directed at normative *realism* in particular, and it is launched from the standpoint of other plausible normative premises. The problem that has been identified is an *internal* incoherence among a set of normative judgments that Dworkin and many others accept: on the one hand, Dworkin accepts that there are independent normative truths, and on the other he thinks that he is not hopeless at discerning what they are, but it turns out that these two normative views are incompatible when combined with other normative assumptions he also would accept—in particular, reasoning about the odds that among all the possible coherent normative systems, one’s own is the right one, coupled with the observation that one has no non-trivially-question-begging evidence that one’s own system is the right one. Nor am I assuming that the *answer* to this challenge needs to be external. On the contrary, I am assuming that the reply, if there is

²⁵ Of course none of us is actually internally consistent. But we may assume that we are for the sake of argument, so as to focus our attention on the important point.

one, will appeal to further normative premises, and that this is fully legitimate; I agree with Dworkin that one can reach no conclusion about the realism/antirealism debate without relying on further normative premises. What form might such an internal reply take? It's here that thinking about the role of causal forces in shaping our systems of normative judgment comes to the fore, and where it will help to turn to the next objection.

11. Second objection: The argument proves too much

According to a second objection, the problem with the argument I have offered is that it proves too much. I am purporting to raise a skeptical challenge that is specific to the normative domain—a challenge that undermines our belief in independent normative facts and our ability to recognize them without similarly undermining our belief in, for example, independent facts about our manifest surroundings and our ability to recognize them. But one might think that the argument I have offered is nothing more than a perfectly general skeptical worry deployed selectively.²⁶ After all, surely there are innumerable equally coherent views about our manifest surroundings; why think we won this “lottery” either? One might argue that here too, just as in the normative case, we will ultimately be able to give only trivially question-begging reasons for thinking we're not hopeless, and hence no special problem has been raised about the normative domain.

This is not so, however; the problem I have raised is specific to the normative domain. To see this, consider the analogous challenge regarding our views about our manifest surroundings. There are, I agree, innumerable equally coherent views about one's manifest surroundings, and causal forces have undoubtedly shaped which of these views one has ended up arriving at. So why think, in this case, that the causes would have led us to views that are roughly correct, when there are countless other, equally coherent views to which the causes could have led us? We have already canvassed the basic answer to this question in section 7. The basic answer is that features of our manifest surroundings have causal powers—in particular being the kinds of things that can kill us, injure us, help us, and so on—and because

²⁶ Cf. Nicholas Sturgeon's criticism of Gilbert Harman's argument in Sturgeon's “Moral Explanations.”

of this, they are the kinds of things that it would have promoted the survival and reproduction of our ancestors to track with reasonable accuracy. The same cannot be said of normative reasons. Normative reasons do not have causal powers—or at least so we are assuming with Dworkin and the non-naturalists in this paper—and because they don't, we are unable to explain how an ability to recognize them would have promoted reproductive success and thus been selected for. Thus, when we survey the countless possible coherent systems of belief regarding our manifest surroundings, we have good reason to think that causal forces would have led us to be able to hit upon roughly accurate ones, whereas when we survey the countless possible coherent systems of normative belief, we have no reason to think this.

One might not be satisfied with this, however. One might think that I am still unfairly holding normative facts to a standard to which I am not holding facts about our manifest surroundings. In particular, one might complain that while I have indeed offered a reason to think that causal forces led us to hit upon roughly accurate views about our manifest surroundings, the reason I have offered is question-begging. After all, it appeals to evolutionary theory, which in turn of course depends for its justification on all kinds of observations about our manifest surroundings. So in answering the skeptical challenge in this case, I have merely assumed that we are not hopeless at recognizing facts about our manifest surroundings, thus begging the very question at issue. According to this objection, if we apply the same standard in the normative case—permitting the assumption that we are not hopeless at recognizing facts about reasons and value—then we can answer the skeptical challenge in that case as well. For example, if we are permitted to assume that things such as family and friendship are valuable, then it will be no mystery, from an evolutionary point of view, why causal forces led us to these true normative views (the rough idea being that it tended to promote our ancestors' reproductive success to value their family members and friends).

I have no dispute with two points being made here. First, I agree that the answer to the skeptical challenge I have offered in the manifest surroundings case is ultimately question-begging—ultimately merely assuming, in its appeal to evolutionary theory, that we are not hopeless at recognizing facts about our manifest surroundings. Second, I agree that in seeking to answer the skeptical challenge in the

normative case, the normative realist is equally well entitled to offer an answer that is ultimately question-begging—an answer that ultimately merely assumes that we are not hopeless at recognizing normative facts. Both of these points, after all, are fully consistent with the kind of skeptical challenge I intend to be posing: an *internal* skeptical challenge whose answer, I am assuming, will also be internal. I am by no means requiring that we explain—from a standpoint that makes no assumptions whatsoever about our manifest surroundings—why we are not hopeless at recognizing objects in our manifest surroundings, and no more so am I requiring that we explain—from a standpoint that makes no normative assumptions whatsoever—why we are not hopeless at recognizing our reasons.

What then is the difference between the two cases? The answer lies in the distinction between answers to the skeptical challenge that provide *internal reason* to think that the causes might have led us to the truth versus answers that provide *no reason whatsoever* to think that the causes might have led us to the truth.²⁷ The general question we are asking in both the manifest surroundings case and the normative case is “Why think that the causes described by our best scientific explanations would have led us to the truth in this domain?” In answer to this question, it is unsatisfactory to reply, “My judgments in this domain are true, and they’re also the ones that the causes described by our best scientific explanations led me to.” Such a reply offers no *reason* for thinking that the causes led us to the truth; it merely reasserts that they did. So, for example, in the manifest surroundings case, it will not do to argue “There are objects *X*, *Y*, and *Z* in my manifest surroundings and that’s also what the causes led me to believe.” Similarly, in the normative case, it will not do to argue “*X*, *Y*, and *Z* are valuable and that’s also what the causes led me to believe.” The problem for normative realism is that in the normative case, empty replies of this sort are all that can ever be given (assuming one is a realist), whereas in the manifest surroundings case, a more informative reply is available.

More specifically, in the manifest surroundings case, we needn’t make any immediate, substantive assumptions about the objects in our manifest surroundings in order to answer the skeptical

²⁷ For related discussion, see chapter 13 of Gibbard’s *Thinking How to Live*, where Gibbard makes a similar point but draws a different ultimate conclusion. I am indebted to Gibbard’s discussion here.

challenge. I needn't assume, for example, that there is a table in front of me or a chair behind me.

Instead, I can make a claim about the general nature of objects in one's manifest surroundings: they are things with causal powers. And since things with causal powers in one's immediate surroundings are the kinds of things that can affect one's survival and reproduction, we can see why selective pressures might well have led us to form accurate beliefs about them. There is no guarantee of this, of course—this is no answer to the *external* skeptic about our ability to perceive objects in our manifest surroundings—but we nevertheless have good *internal* reason to think we are reliable in this domain.

In the normative case, however, we can say no such thing. One of course cannot run into, or eat or be eaten by, reasons and values. More generally, if we adopt the non-naturalist realist's conception of normative truth—as independent and yet lacking in causal powers—there is no reason to think that natural selection, or for that matter any other causal process, would shape us in such a way that we would be able to track such truths. The objector argues that family and friendship (say) are valuable, and explains why evolutionary forces would have led us to think this. But no matter how we dress it up, this kind of account is ultimately just of the unacceptable form “X, Y, and Z are valuable, and that's also what the causes led us to believe.” It provides no *reason* to think that the causal forces described by our best scientific explanations shaped our normative judgments in ways that might have led those judgments to track the truth; it merely confidently reasserts that they did. To put it another way: One may explain each *side* of the coincidence in as much depth as one likes—going into normative depth about why family and friendship are valuable, and scientific depth about why we were selected to think this. But all such depth goes nowhere toward explaining the thing that really needs to be explained, namely the *coincidence itself*.

Are there other things the realist might say to try to explain why causes would have worked on our normative judgments in such a way as to enable us to track independent truths about reasons and values? We talk about the causal powers of objects in our manifest surroundings; what comparably general things might we say about reasons and values? We might say, for instance, that recognizing them involves recognizing how one should spend one's time, how one should treat friends, family, strangers, other animals, and so on. This is true, of course, but it tells us nothing about why causal forces would

have led us to be able to recognize them. What if we said that reasons have to do with, among other things, the welfare of ourselves and others and how best to promote it, and so it's no wonder if causal influences shaped us to be able to recognize them? The problem with this reply is that *welfare* is a normative concept, and so in making this reply either one is or one isn't making assumptions about what constitutes a creature's welfare. If one is *not* making any assumptions about what constitutes a creature's welfare (leaving it completely open whether it consists in survival or lack thereof, for example, or intact bones or broken ones), then it is no more clear why causes would have shaped us to track our welfare than it is why they would have shaped us to track reasons and values. If, on the other hand, one *is* making assumptions about what constitutes a creature's welfare—assuming, for example, that other things being equal, survival, security, a long life, lack of broken bones, and so on are all part of a creature's welfare—then one has merely lapsed back into offering a reply of the unsatisfactory form, “X, Y, and Z are valuable (or good, or what we have reason to do, or constitutive of welfare, etc.) and that's also what the causes led me to believe.” Again one is merely insisting that the causes led us to the truth without providing any reason to think so.

We may now see how the exact same skeptical challenge is being issued in both the manifest surroundings case and the normative case, and why the challenge succeeds in the latter case (assuming one is a normative realist) but not the former. In both cases, we (1) consider the nature of truth *as such* in the given domain—meaning that we think about the general nature of the subject matter without making any substantive assumptions about what in particular is true of that subject matter; (2) examine our best going scientific account of the causes that shaped our judgments regarding that subject matter; and then (3) ask whether there is any (internal) reason to expect that the causes in question would have led us to be able to track the type of truth in question, given our understanding of its nature. In the manifest surroundings case, as I have argued, there is such a reason, whereas in the normative case—assuming the truth of non-naturalist realism—there is not. The larger upshot is this. In both cases, there are, in principle, countless equally coherent systems of belief available to us. In the manifest surroundings case, since we have some reason to expect that the causes in question would have led us to the truth, we are

able to see ourselves (from an internal point of view) as having arrived at the system we did *non-arbitrarily*. We may therefore retain reasonable confidence that we are reliable in this domain. In the normative case, in contrast, since we have no reason to expect that the causes in question would have led us to the truth, we are forced to view ourselves (again, from an internal point of view) as having arrived at the system we did *arbitrarily*. This is where the causes landed us, in other words, but at bottom we have no reason to think they had anything to do with the truth. In effect, we must view ourselves as having selected our system of normative judgments at random—by way of pure causal contingency—from among a countless number of equally coherent normative systems. It's exactly like the New York State lottery case, then, when you have no reason to think you won: given the odds against it, you'd better conclude that in all likelihood you didn't.

To drive home the point that no double standard is being employed, notice finally that there is no reason in principle why the skeptical challenge couldn't be answerable in the normative case but not in the manifest surroundings case. We need only vary the conceptions of truth or the nature of the scientific explanations in play. For example, suppose that in the manifest surroundings case, it turns out that our best scientific understanding of the causes that led us to our judgments turns out to be not an evolutionary account, but rather the view that we have been hypnotized into our views by a whimsical hypnotist who is just out to entertain himself. In such a case, we would lose our internal reason to think the causes would have led us to hit on the truth about our manifest surroundings, and so, given the sheer odds against landing on the right coherent system randomly (the hypnotist's whimsy being as good as random with respect to the truth), we would be forced to conclude that in all likelihood we were hopeless at discovering the truth about our manifest surroundings. And it's worth noting that here too, just as in the normative case, it is clearly unacceptable to try to answer the skeptical challenge merely by stating that objects *X*, *Y*, and *Z* are in our manifest surroundings, and then explaining in detail the technique the whimsical hypnotist used to get us to believe this.

Just as we could conceivably *fail* to meet the skeptical challenge in the manifest surroundings case, we could conceivably *meet* it in the normative case. Indeed, I have already sketched out how this is

possible in section 6. In particular, if we vary our conception of the truth as such—rejecting the non-naturalist realist’s conception and adopting a constructivist conception—we are able to offer good internal reason to think we’re not hopeless at tracking the normative truth. For suppose, in accordance with my version of constructivism, that the fact that *X* is a reason for agent *A* to *Y* is constituted by the fact that the judgment that *X* is a reason (for *A*) to *Y* withstands scrutiny from the standpoint of *A*’s other judgments about reasons. While of course more details of the view need to be filled in before a full assessment can take place, I hope it is clear enough in outline how we are able to track facts about what withstands scrutiny from various normative standpoints, including our own, and that this is unmysterious from a naturalistic point of view.²⁸ On the constructivist conception of truth, then, we have internal reason to think we are reliable judges of normative truth.

Thus, as I’ll explain in more depth below, I believe we are forced to move to a constructivist conception of normative truth in order to avoid the skeptical result that issues from realism. But this move might seem a bit odd. Is it legitimate to change one’s conception of truth in a given domain so that one may view oneself as reliable? Suppose our best evidence suggested that a whimsical hypnotist was behind our beliefs about our manifest surroundings. Surely it wouldn’t make sense to alter our conception of truth about our manifest surroundings so that we’d no longer have reason to think we were hopeless—altering it, for example, to the view that the truth about objects in our manifest surroundings consists in whatever strikes the fancy of the hypnotist. I agree this wouldn’t make sense. But I think that every domain of judgment needs to be looked at individually, and—as Dworkin would agree—it’s a question of what’s most plausible all things considered. In my view, what’s most plausible all things considered in the normative case is to change our conception of normative truth—to reject realism and accept constructivism (assuming we weren’t constructivists already). So let me now turn to doubts about that claim of overall plausibility.

²⁸ For further discussion, see “Constructivism About Reasons.”

12. *Third objection: The overall balance of plausibility*

Assume that my replies to the previous two objections are satisfactory. The defender of non-naturalist realism still has another objection he can raise. Dworkin presents it in the following passage:

Let us accept, for the sake of argument, that we are forced to choose between the following two propositions. (1) Human beings have a special though sometimes fallible faculty of judgment that enables us to decide which moral claims to accept or reject, a capacity whose malfunctioning may sometimes result only in moral misjudgment with no spillover impairment of other cognitive activity. (2) There is no moral objection to exterminating an ethnic group or enslaving a race or torturing a young child, just for fun, in front of its captive mother. Which should we abandon? (pp. 117-118)

According to Dworkin, here as in all cases we must abandon the proposition that seems least plausible all things considered. “If forced to choose between the two propositions just described,” Dworkin says, “I would accept the first and reject the second,” and he believes the rest of us would too. The argument here is that although the skeptic may raise some reasonable worries about the epistemology associated with normative realism, realism—even accompanied by these worries—is *still* more plausible than the alternative; the overwhelming plausibility of realism when compared with antirealism is enough to rebut skeptical attacks of the type I have raised.²⁹

In reply, let me begin by granting three things. First, I agree with the important point being made here that ultimately, the realism/antirealism debate must be decided on holistic considerations. The mere fact that realism brings with it some apparent problems—even very serious ones—is not enough to defeat the view if the alternative, antirealism, is even more implausible in comparison. Second, I view the general form of Dworkin’s argument here as entirely legitimate. Since we are understanding the realism/antirealism debate as a normative debate like any other—just taking place at a higher level of abstraction than usual—ultimately the argumentative situation here is no different from any other case in which we find two or more of our normative convictions in tension—for instance when we notice that the principle “One should always save the greatest number” is in tension with the view that “One should not sacrifice an innocent to harvest his organs for use by others.” Just as in this organ case we can do nothing

²⁹ [Cite Moore.]

but opt for that conviction which, on reflection, seems most plausible all things considered, so I agree we can do nothing in the realism/antirealism debate but opt for the conviction which, on reflection, seems most plausible all things considered. Third and finally, I agree that if the choice is as Dworkin describes it—between realism coupled with a somewhat mysterious normative epistemology and antirealism coupled with the view that there is no moral objection to torturing a child in front of its captive mother, and so on—then the former view *is* more plausible than the latter, and we are justified in cleaving to realism.

This, however, is not the choice we face. Instead, realism commits us to something much worse than Dworkin's proposition (1), whereas antirealism commits us to something much less counterintuitive than Dworkin's proposition (2). The real choice, as I see it, is between the following two propositions: (1') In all likelihood I am hopeless at recognizing my reasons; and (2') Some conceivable agents have reason to exterminate an ethnic group or enslave a race or torture a young child for fun in front of its captive mother, but most real life human beings have no such reasons, and if we ever encounter any who do, then we (you and I and the vast majority of human beings) have reason to band together against them, lock them up, and throw away the key.³⁰ Once we see that this is the real choice, it becomes clear that holistic considerations of the same kind in play in any normative debate force us to antirealism.

To see that this is the real choice, consider first (2') versus (2). Dworkin's proposition (2)—according to which “There is no moral objection to exterminating an ethnic group or enslaving a race or torturing a young child, just for fun, in front of its captive mother” is a deeply misleading, boogeyman

³⁰ Note that in stating proposition (2'), I am assuming that *Humean* antirealism is true. The argument becomes even easier if Kantian antirealism is true. Then, rather than having to concede that some agents have reason to torture young children, etc., we can insist—as for instance Korsgaard would—that *no* agents have such reasons, so the choice is really between (1') In all likelihood I am hopeless at recognizing my reasons; and (2'') No one has reason to exterminate an ethnic group or enslave a race or torture a young child for fun in front of its captive mother; rather, everyone has reason to treat others as ends in themselves. Here the choice becomes a no-brainer in favor of antirealism. However, since I doubt that Kantian versions of antirealism can be made to work, I focus in this section on the more difficult case of the choice between realism and Humean antirealism, arguing that Humean antirealism still wins out.

characterization of what is involved in accepting antirealism. A sophisticated antirealist will agree to no such thing.

First, Dworkin's proposition (2) is couched not in terms of normative reasons in general, but in terms of what there is or isn't "moral objection to." Straight off, this is unfair to the antirealist, for one might think that it's a conceptual truth or near conceptual truth that there is "some moral objection" to torturing a child for fun—such that any "morality" that denied this wouldn't be recognizable as a brand of *morality* at all. Because it denies such a fundamental moral platitude, proposition (2) sounds almost crazy. But morality/reasons externalism, according to which we may not always have *reason* to do what we morally ought to do, is a well-known option that the antirealist can adopt regarding such a case.³¹ In other words, the antirealist might grant that it is *morally objectionable* to torture a child for fun, but then deny that all agents always have reason to be moral. To avoid prejudicing the debate against the antirealist with an implicit appeal to platitudes associated with the concept of *morality*, then, the claim in question should be couched in terms of reasons. Of course it still sounds deeply implausible to say that "There is no reason not to exterminate an ethnic group or enslave a race or torture a young child for fun in front of its captive mother," but this brings us to a second point about the formulation of proposition (2).

The second point is this. An antirealist will not accept blanket statements of the form "There is no reason not to *Y*" or "There is reason to *Z*." The antirealist will insist that at least when we're being careful and not leaving things implicit, we must always specify *for whom* a given consideration is or is not a reason; there is no such thing as something's being a reason or not full stop—reasons are always reasons *for* some agent.³² Thus, no antirealist worth her salt would assent to "There is no reason not to exterminate an ethnic group or enslave a race or torture a young child for fun." Instead, the most the antirealist will agree to is that "Some agents, because they possess a starting set of evaluative attitudes

³¹ The label *morality/reasons externalism* is Stephen Darwall's in his taxonomy of varieties of internalism and externalism in "Reasons, Motives, and the Demands of Morality: An Introduction," in Stephen Darwall, Allan Gibbard, and Peter Railton, *Moral Discourse and Practice* (New York: Oxford University Press, 1997), pp. 305-312.

³² I discuss this point at more length in "Constructivism About Reasons."

very different from our own, have no normative reason not to exterminate an ethnic group, enslave a race, or torture a young child for fun. An ideally coherent Caligula would be an example of such an agent.”

And note moreover that *some* antirealists—ambitious *Kantian* antirealists such as Korsgaard—would not even assent to this much, claiming instead that it’s entailed from within every practical point of view that one has reason not to do such things. If such Kantian antirealists can make their arguments work, then they can reject Dworkin’s proposition (2) out of hand. Until it has been shown that such arguments do not work, then, it is entirely premature to attribute a claim anything like (2) to the antirealist.

What if one is a Humean antirealist, though? This question is important, since I, and I think many others, are pessimistic about the prospects of making Kantian antirealism work. If one is a Humean antirealist, then one must admit that *some* agents will have no reason not to torture children for fun, and so on. But this claim too is much more plausible than it might seem once properly understood. Indeed, once Humean antirealism is fully spelled out, I think it becomes clear that the view allows us to say almost everything we ever were inclined to say about people’s reasons, with the exception of a swath of extremely strong realist claims. Here there are at least five points to be made, which I will have to pass over very quickly.³³

First, it is important to note that on a Humean antirealist view, in order for an agent to have no reason not to torture children for fun, strict conditions will have to be met; roughly, it will have to be the case that an ideally coherent and fully informed version of herself would judge that she has no reason not to torture children for fun. And given human nature as it actually is, it’s the very rare real life person who would ever meet these conditions. Second, if we imagine in full, vivid detail what such a person would have to look like, the view that such a person *would* have reason to torture others becomes intuitively much more plausible. Third, since no real life human being is anywhere close to ideally coherent, in speaking to others who disagree with us on matters such as torture, we should always proceed on the assumption that there *is* a shared fund of evaluative starting points such that it follows from our

³³ I say more about these points in “In Defense of Future Tuesday Indifference” (draft).

interlocutor's own judgments too that he or she never would have reason to torture a child for fun.

Fourth, as Bernard Williams notes in his defense of reasons internalism, even if it's true that a person has *reasons* that are repugnant to us, there are still plenty of other things we can say about this person³⁴—for instance that he is despicable to us, and that it's sad, from our point of view, that he and his reasons are like this. And finally, it must be emphasized that when and if we ever encountered such a person, the rest of us would have every reason to lock him up, defend ourselves against him, and to try to change him if we could. Our dealings with such a person would be like our dealings with an intelligent alien species intent on eating us: we might grant that members of this species were *justified* in trying to eat us, but we would nevertheless of course fight them.

So much for Dworkin's proposition (2), then: a sophisticated antirealism is much less counterintuitive than this proposition suggests. Antirealism and its intuitive costs are much better summarized by a proposition along the lines of (2')—and that's on the assumption that Kantian versions of antirealism fail, something that for all we have said here, might turn out to be false, rendering antirealism even more fully in accord with our pre-theoretical intuitions.

Now consider (1') versus (1). According to Dworkin's proposition (1), realism brings with it an admittedly somewhat mysterious epistemology: we must believe that we have a special but fallible faculty that enables us to grasp independent normative truths, but whose malfunctioning may sometimes result only in normative misjudgment without any "spillover impairment of other cognitive activity" (p. 117). This, while perhaps not entirely satisfactory, Dworkin seems to concede, is intuitively much more plausible than proposition (2). We have now seen that there is no need to accept (2)—that the much more palatable (2') is an option. And now we must add to that that the situation is actually much worse for the realist than a commitment to (1). If the arguments I'm making in this paper are correct, then the realist is committed not only to a somewhat mysterious epistemology, but also to the much more radical conclusion that—given the phenomenal odds against having won the normative lottery and the utter lack

³⁴ Williams, "Internal and External Reasons," p. 110. [Discuss why I think the adjectives Williams uses are problematic, though.]

of any (internal) reason to think we did—we are in all likelihood hopelessly off track with our normative judgments. And it's not just that this conclusion is extraordinarily implausible, which it is (surely each of us thinks that we're at least *okay* at figuring out what many of our reasons are). It's also that one *must* reject this conclusion if one is to go on making normative judgments at all.³⁵ It is incoherent to affirm, for example, that "I should do *Y*, but in all likelihood I'm completely hopeless at recognizing what I should do." If one genuinely thinks one is in all likelihood hopeless at recognizing one's reasons, one will be paralyzed and unable to proceed with normative reasoning; and if one does go ahead with one's practical reasoning, then one is, in doing so, implicitly holding the view that one is not in all likelihood completely hopeless.

The overall balance of plausibility, then, is quite different from what Dworkin suggests. Not only is antirealism much more palatable once correctly understood, but the alternative—realism—commits us to a conclusion that is not only extremely implausible—much more implausible, it seems to me, than the idea that some agents could have reason to torture others for fun—but also literally unacceptable, being incompatible with going forward as a practical reasoner at all. Dworkin writes, "If you can't help believing something, steadily and wholeheartedly, then you'd better believe it" (p. 118). I agree with this. But the thing I claim we can't help believing, steadily and wholeheartedly, is that we have some idea of what our reasons are. And this conviction is what drives us to antirealism.

13. Fourth objection: Isn't antirealism self-defeating?

Antirealism, as I myself have granted, is a substantive normative view. Yet it is also an extremely general one, according to which there are no normative truths that hold independently of agents' evaluative attitudes and what follows from within the standpoint constituted by those attitudes. This, however, raises the worry that antirealism could apply to itself and be self-defeating. To consider this objection, it will be useful to focus on the specific antirealist view I favor. As we saw earlier, according to this constructivist view, the fact that *X* is a reason for agent *A* to *Y* is constituted by the fact

³⁵ [Cite Gibbard, *Thinking How to Live*.]

that the judgment that *X* is a reason (for *A*) to *Y* withstands scrutiny from the standpoint of *A*'s other judgments about reasons. One might argue that this view is self-defeating in the following way.³⁶

Suppose that Ann the realist makes the following normative judgment (call it *J*): “*X* is a reason for me to *Y* and it would be such a reason even if I judged that it wasn't and even if that very judgment of mine withstood scrutiny from the standpoint of all of my other judgments about reasons.” This is just a straight affirmation of realism in opposition to constructivism. Suppose, moreover, that *J*, which is itself a normative judgment, withstands scrutiny from the standpoint of all of Ann's other normative judgments. Then there seems to be a problem. On the one hand, *J* would seem to be false on my view, since it's in direct contradiction to my constructivist proposal. On the other hand, *J* would seem to be true on my view, since it stands up to scrutiny from the standpoint of Ann's other normative commitments.

My reply to this is to deny that *J* stands up to scrutiny from the standpoint of Ann's other normative commitments. Contrary to what the example stipulates, *J* does not withstand such scrutiny. My claim is that no matter what starting set of normative judgments one begins with—no matter how causes shaped one's initial set of normative commitments—it follows from those commitments that constructivism is true. In other words, constructivism withstands scrutiny, and realism does not, from the standpoint of *every agent's* normative commitments, including Ann's, even though she does not recognize this. The argument for this is just the argument of this paper. Every agent is capable of viewing herself from both the practical and the theoretical point of view. But this means that the practical/theoretical puzzle arises for her. And no matter what her set of normative judgments, it will follow from those judgments that realism is an unacceptable solution to this puzzle. This is because there's one thing that every agent is committed to in virtue of making any normative judgments at all, namely that she is not hopeless at recognizing reasons. But *realism*, as I have argued, has as its consequence that in all likelihood she *is* hopeless, since the agent has no way of seeing why the causal forces which she knows to have shaped her normative judgments would ever have led her to true

³⁶ I am indebted to Ronald Dworkin for raising this objection.

normative judgments. Since it is unacceptable to think she is hopeless, she must reject realism. All this follows from within her own point of view no matter what set of normative judgments she begins with.

Constructivism is not self-defeating, then, for no matter what one's starting set of normative judgments, constructivism follows from within the standpoint constituted by those judgments. It thus meets its own standard, and no problem arises when we apply it to itself. I said earlier (in section 2) that my version of antirealism is—with an important qualification—Humean. It should now be clear what that qualification is. While my version of antirealism is Humean in the sense that I am generally skeptical of attempts to show that strong, substantive conclusions about our reasons are entailed from within every practical point of view—believing that attempts to derive morality from within every practical point of view fail, for example—there is one strong, substantive conclusion about our reasons that I think *is* entailed from within every practical point of view, namely constructivism itself. Thus, although my version of antirealism is largely Humean—insisting that agents' reasons vary radically depending on their starting set of normative commitments—it is Kantian at the very bottom, holding that this very truth—about the radical dependence of reasons on our starting set of normative commitments—follows from *any* starting set of normative commitments.

Since the constructivist can see her way to constructivism no matter what starting set of normative judgments she accepts or might have accepted, she needn't view herself as merely "lucky" in having reached this conclusion. She needn't view herself as among a mysteriously privileged "normative elect." She sees that her constructivist conclusion is not contingent on her particular starting set of normative convictions and the causes that gave rise to them. Unlike the realist, she has reasons all the way down. In this way, she finds harmony between the practical and theoretical points of view.

14. Conclusion

It's important to emphasize that on a constructivist account, there is still normative truth and objectivity. It's just that the level of objectivity isn't quite as robust as the realist thought. While it's entirely possible to make normative errors on a constructivist view, and we do so all the time, the

standards of correctness that determine what are and aren't errors are *ultimately* set by the normative judgments of the person whose reasons are in question. This is the way in which we're forced to rethink normative truth and objectivity.

Toward the end of his article, as he finishes his case against external forms of normative skepticism, Dworkin remarks that the "only kind of skepticism that counts, anyway, is the really disturbing kind, the chilling internal skepticism that grips us in a dark night, when we suddenly cannot help thinking that human lives signify nothing, that nothing we do can matter when we and our whole world will in any case perish in a cosmic instant or two" (p. 129). Constructivism is a form of internal skepticism, but it's not of the chilling kind. On the contrary, if constructivism is right, then there is a sense in which it's always mistaken to be in the grip of the view that nothing we do can matter. What we do matters so long as we *take* it to matter, for according to constructivism that (with some complications) is just what it *is* for something to matter. The real threat here, then, is merely causal (as opposed to rational): it is that we will somehow cease taking things to matter. But I think that's just the real (and merely causal) threat of death (no less frightening, of course, for being merely causal). So long as we're alive and taking things to matter, though, things *do* matter.³⁷

³⁷ For their helpful comments on an earlier version of this paper, I am indebted to A. J. Julius, Karl Schafer, Nishi Shah, and David Velleman, as well as audience members at Mount Holyoke College. I am also indebted to members of the Science and Religion Project at NYU, especially Paul Boghossian, Matt Evans, Philip Kitcher, Matthew Kotzen, Anna-Sara Malmgren, Thomas Nagel, Alvin Plantinga, Michael Strevens, and Roger White.