

---

## MORAL SKEPTICISM FOR FOXES

DANIEL STAR\*

During the metaethics panel at the *Justice for Hedgehogs* conference, the audience was presented with something of a spectacle: three moral realists speaking up for, and sometimes passionately defending, moral skepticism – Russ Shafer-Landau, Michael Smith, and myself.<sup>1</sup> The target was Part I of Ronald Dworkin’s *Justice for Hedgehogs*, in which Dworkin attempts to establish that a thoroughgoing “external” moral skepticism is unintelligible.<sup>2</sup> In large part, his attempt to do this is based on arguments that lead to the conclusion that metaethics – the field in which it is normally thought that moral realists and moral skeptics trade arguments – is not a genuine field of philosophical enquiry.<sup>3</sup> A cynic at the conference might have been forgiven for thinking that the only reason a group of moral realists were finding it necessary to defend moral skepticism was because their jobs depend on their field being recognized as legitimate. (Why else would committed realists defend skepticism rather than simply celebrate its demise?) Fortunately, there is a much better explanation available: Dworkin’s discussion of moral skepticism, as interesting as it is at times, is fundamentally unfair to the moral skeptic.

Dworkin’s approach is unfair because it attempts to silence the moral skeptic before she has even presented an argument, using resources that are not strong enough for such a Herculean task. There are, in fact, as many moral skepticisms as there are arguments for moral skepticism, and the devil is always in the details of particular arguments. Rational debates between moral realists and moral skeptics are essentially philosophical debates, and philosophy is essentially a fox-like activity, even if it is also sometimes a

---

\* Assistant Professor of Philosophy, Boston University. Many thanks to Aaron Garrett, our chair at the conference, for a very helpful conversation about a draft of my presentation. I would also like to thank Stephen Kearns for the usual highly stimulating conversations, as well as James Fleming for the invitation to be a speaker at his extremely well-organized conference.

<sup>1</sup> A fact first commented upon by Aaron Garrett. It should be noted that throughout this essay “moral skepticism” is taken to be synonymous with “moral antirealism,” rather than with “uncertainty about morals.” This is a rather confusing, but now standard philosophical usage of the term that Dworkin himself adopts.

<sup>2</sup> RONALD DWORKIN, *JUSTICE FOR HEDGEHOGS* (forthcoming 2010) (Apr. 17, 2009 manuscript at 16-63, on file with the Boston University Law Review).

<sup>3</sup> *Id.* (manuscript at 56).

hedgheg-like activity.<sup>4</sup> Nonetheless, one might reasonably hope that the moral realist can establish that her position is the *default* option – the starting point which the skeptic needs to argue one out of – and if realists manage to establish this much they will, in effect, be saving what might be thought to be the guiding intuition in this part of Dworkin’s book, i.e., the intuition that substantive moral judgments stand in the way of moral skepticism.<sup>5</sup>

---

<sup>4</sup> In the famous essay that introduced the relevant distinction between foxes and hedgehogs, Isaiah Berlin includes Plato amongst the hedgehogs, who are thinkers that “relate everything to a single central vision, one system less or more coherent or articulate, in terms of which they understand, think and feel – a single, universal, organizing principle in terms of which alone all that they are and say has significance.” ISAIAH BERLIN, *THE HEDGEHOG AND THE FOX* 1-2 (1953). Plato is clearly a philosopher if anyone is, even if one does not agree with A.N. Whitehead that the history of Western philosophy is “a series of footnotes to Plato.” ALFRED NORTH WHITEHEAD, *PROCESS AND REALITY* 39 (David Ray Griffin & Donald W. Sherburne eds., corrected ed. 1978). This does not undermine my claim that philosophy is essentially a fox-like activity. Berlin is far from alone in describing Plato in the way that he does but, arguably, Plato was much closer to contemporary philosophers in his methodology than Berlin suggests. As Bernard Williams says,

[I]t is a mistake to suppose that Plato spends his time in the various dialogues adding to or subtracting from his system. Each dialogue is about whatever it is about, and Plato pursues what seems interesting and fruitful in that connection. . . . We should not think of him as constantly keeping his accounts . . . .

BERNARD WILLIAMS, *Plato: The Invention of Philosophy, in THE SENSE OF THE PAST: ESSAYS IN THE HISTORY OF PHILOSOPHY* 148, 154 (Myles Burnyeat ed., 2006). In his original essay, Berlin himself admits that “like all over-simple classifications of this type, the dichotomy becomes, if pressed, artificial, scholastic and ultimately absurd.” BERLIN, *supra*, at 2. It became clear during the conference that Dworkin himself does not wish to make too much of the distinction between foxes and hedgehogs, despite the choice of title for his book.

<sup>5</sup> See DWORKIN, *supra* note 2 (manuscript at 42). There is much evidence that moral realism has already become the default option in metaethics: in particular, the opposition does not now consist of philosophers as opposed to the spirit of moral realism as Ayer and Mackie were – they both claimed that first-order moral statements always lack the property of being true, each in his own way – but rather consists of quasi-realists, ecumenical expressivists, irrealists, fictionalists, etc., who would pretty much all like to find a way of agreeing with the realist that each of us can know (or “know”) that torturing people for fun is wrong, and of explaining how it is that we can make logically compelling inferences that contain moral claims. For instance, expressivists have squarely faced up to the challenge of solving the Frege-Geach problem, and it is apparent that a solution to this problem will need be very complex. See generally MARK SCHROEDER, *BEING FOR: EVALUATING THE SEMANTIC PROGRAM OF EXPRESSIVISM* (2008). It is at this point, but not necessarily elsewhere, that the traditional realist has a simpler and more elegant story to tell: we can know that torturing people for fun is wrong because moral judgments are beliefs, and some such beliefs are true and warranted, epistemically speaking (in at least one of the standard ways that beliefs can be warranted). It is not at all surprising that moral realism is, in this way at least, the default option in contemporary metaethics if the “reasons as evidence” thesis I describe below is correct, for this thesis may be thought to provide an explanation as to why moral realism is

I agree with Dworkin that our substantive moral judgments do provide an important barrier to skepticism – a strong fence to keep out foxes, if you like – and I will attempt to explain here how exactly they do this, in a way that does not rule out all forms of moral skepticism *a priori*.<sup>6</sup> If my attempt to do this is successful, I will have found a more promising, yet more modest way of protecting moral realism from attack than the one that Dworkin himself offers us in Part I of *Justice for Hedgehogs*.<sup>7</sup>

Perhaps the best place to begin is with Dworkin’s central contention that “skeptical conclusions must themselves be understood as substantive moral opinions – otherwise they make no sense at all.”<sup>8</sup> Consider the following four types of claims: (1) I ought to  $\phi$ ; (2) I ought not  $\phi$ ; (3) it is not the case that I ought to  $\phi$ ; and (4) it is not the case that I ought not  $\phi$ . Taken at face value, Dworkin’s central contention is false, since claims of type (1) or (2) are *logically distinct* from claims of type (3) or (4) – as all four types of claims have logically distinct forms – and the moral skeptic standardly *never* appeals to the first two types of claims, which are the only forms that are *necessarily* substantive, but appeals instead to the second two types of claims, which are *not necessarily* substantive.<sup>9</sup> This is just to say that there is a difference between a first-order moral claim and a second-order claim about a moral claim, which is not itself a moral claim. It is possible to make this even easier

---

the default option. I will ignore quasi-realism in this essay; when I speak of realism and the “reasons as evidence” thesis below, I can be read as endorsing either traditional moral realism, which I favor, or a quasi-realism that succeeds in endorsing the same substantive moral and epistemic claims that traditional moral realism attempts to – assuming it can succeed in doing this. This might be termed “weak” moral realism, but it is a realism strong enough for my purposes here.

<sup>6</sup> There may be a way of ruling moral skepticism out *a priori*, for all that I say here; it is simply not my present goal to discover or promote one, and it seems clear, in any case, that Dworkin does not provide us with one.

<sup>7</sup> DWORKIN, *supra* note 2 (manuscript at 16-63).

<sup>8</sup> *Id.* (manuscript at 18).

<sup>9</sup> I say “not necessarily” because a claim of this second kind can, in some but far from all contexts, take the form of a substantive moral claim. For the sake of illustration, compare “It is not the case that I should donate money to that charity, because there are better charities for me to donate money to” with “It is not the case that I should donate money to that charity, because there are no ‘should’ truths of any kind.” This parallels the difference between “It is not the case that she is a witch, because this spell belongs to someone else” and “It is not the case that she is a witch, because witches do not exist.” The fact that a claim of some kind can in some contexts be entailed by a substantive moral claim is *not* enough to make all claims of that kind substantive moral claims. People can claim that witches (or moral truths) do not exist without being committed to the existence of witches (or moral truths). Dworkin himself made this mistake in his response to our panel at the conference: to my objection below concerning the “ought implies can” principle, he responded that some religious people in the past have explicitly rejected “ought implies can” on substantive moral grounds, and that *therefore* “ought implies can” is itself (always) a substantive moral claim. This is simply a non sequitur.

to see. Consider the moral claim, “I ought to buy this.” Now consider the claim, “The sentence ‘I ought to buy this’ has five words in it.” This second claim is surely not a substantive moral claim, even though the first claim may be.<sup>10</sup> Similarly, garden-variety moral skeptics will say it is always not the case that one ought to  $\phi$ , for each and every possible  $\phi$ , and in saying this they will not be making any substantive moral claims.

Dworkin goes on to say: “We need a more careful statement of external error skepticism: it holds that all *positive* moral judgments are false.”<sup>11</sup> He indicates that by “positive moral judgments” he means only first-order judgments of the form “I ought to  $\phi$ ” and “I ought not  $\phi$ ” and the like. He also mentions judgments that we might make about persons or situations being good or bad, and people possessing virtues or vices.<sup>12</sup> He thus means to exclude judgments that take the form, “It is permissible for me to  $\phi$ .” He then argues that the error skeptic’s thesis must itself be a first-order moral claim, because to say that it is not the case that an act is required or forbidden is to say that it is morally permissible, and moral permissibility claims are also first-order moral claims.<sup>13</sup> Thus, “we can sensibly understand the supposedly second-order philosophical claims of the external skeptics . . . as themselves what they [the skeptics] call substantive, first-order moral judgments.”<sup>14</sup>

Unfortunately, Dworkin fails to appreciate that a moral skeptic will wish to deny that moral permissibility statements are true, just as much as she will wish to deny that statements about what morality requires of us are true.<sup>15</sup> The moral skeptic that Dworkin himself discusses most of all at this point in his text is J.L. Mackie.<sup>16</sup> Mackie is usually interpreted as arguing that all moral

---

<sup>10</sup> In fact, I do not think the first claim is necessarily a moral claim, because it is not the case that “ought” is always used to express moral claims, even when we restrict our attention to normative uses of the word – there are also non-normative uses, such as in “it ought to stop raining today” – but I will ignore this complication, for the most part, following Dworkin’s own concentration on morality, rather than the normative and evaluative spheres in general.

<sup>11</sup> DWORKIN, *supra* note 2 (manuscript at 42). “External error skepticism” is Dworkin’s term for a kind of moral skepticism that sees morality as a whole as a mistaken enterprise, from a position purportedly outside of morality itself. *Id.* (manuscript at 23). An example of *internal* error skepticism would be the kind of skepticism many of us exemplify in relation to traditional sexual morality, where we use some of our substantive moral commitments to criticize particular traditional moral claims. This is obviously not the kind of skepticism that is at issue between Dworkin and opponents such as myself, Michael Smith, and Russ Shafer-Landau. *See generally* Russ Shafer-Landau, *The Possibility of Metaethics*, 90 B.U. L. REV. 479 (2010); Michael Smith, *Dworkin on External Skepticism*, 90 B.U. L. REV. 509 (2010).

<sup>12</sup> DWORKIN, *supra* note 2 (manuscript at 42).

<sup>13</sup> *See id.* (manuscript at 42-43).

<sup>14</sup> *Id.* (manuscript at 42).

<sup>15</sup> *See* Smith, *supra* note 11, at 512.

<sup>16</sup> *See generally* J.L. MACKIE, *ETHICS: INVENTING RIGHT AND WRONG* (1977).

judgments and claims are false, and not just *positive* moral judgments and claims, as Dworkin assumes. To be fair to Dworkin, it is very tempting to think that “it is permissible for me to  $\phi$ ” logically follows from “it is not the case that I ought not  $\phi$ ” – although this is to place more faith in standard deontic logic than a skeptic might wish to (people on all sides accept that rejecting standard deontic logic is nowhere near as outlandish as rejecting logic *per se*). At most, this gives the skeptic a reason to prefer the more contemporary version of moral error theory defended by Richard Joyce.<sup>17</sup> On this type of view, first-order moral claims regarding what is required, forbidden, or permissible are all untrue, but this does not mean that they are false; rather, they altogether lack truth-values (yet still express beliefs when made as sincere claims; beliefs that should also be thought of as lacking truth-values), due to the fact that the presuppositions at play in the moral enterprise – such as that the demands of morality are categorical (that is, apply to all of us regardless of any particular desires we happen to possess) – are false.<sup>18</sup>

Dworkin fails to show that the claim that all positive moral claims are not true is itself a substantive moral claim, and he more generally fails to show that all second-order claims about morality are (also) first-order substantive moral claims. One should look for better ideas in Dworkin’s text. Perhaps the more fundamental and interesting claim of Part I of *Justice for Hedgehogs* is that Hume’s principle – the very plausible principle that it is impossible to logically derive an *ought* from an *is* – rules out external moral skepticism, because it makes it *impossible* to defend.<sup>19</sup> Dworkin himself refers to Hume’s principle as “the anthem of Part I.”<sup>20</sup> If Dworkin can show that every argument for – rather than every claim expressing – moral skepticism must contain a previously overlooked substantive moral premise, then he may succeed in treading in Hume’s footsteps. Hume pointed out that people often leap from purely descriptive claims to normative claims, without noticing,

as this *ought*, or *ought not*, expresses some new relation or affirmation, ’tis necessary that it shou’d be observ’d and explain’d; and at the same time that a reason should be given, for what seems altogether

---

<sup>17</sup> See generally RICHARD JOYCE, *THE MYTH OF MORALITY* (2001). For more details on Richard Joyce and his moral error theory, see Shafer-Landau, *supra* note 11, at 489.

<sup>18</sup> Sometimes Mackie makes statements that can be read as suggesting that he was drawn to such a view himself. He writes: “[T]he denial of objective values will have to be put forward . . . as an ‘error theory’, a theory that although most people in making moral judgments implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false.” MACKIE, *supra* note 16, at 35. “These claims” might refer either to the first-order moral judgments themselves, or to the constantly repeated implicit claims to objective prescriptivity. The second reading would give us Joyce’s type of error theory.

<sup>19</sup> DWORKIN, *supra* note 2 (manuscript at 43).

<sup>20</sup> *Id.* (manuscript at 15).

inconceivable, how this new relation can be a deduction from others, which are entirely different from it.<sup>21</sup>

Unfortunately, Dworkin nowhere makes good on his goal. This is hardly surprising, since there are *many* arguments for moral skepticism in the literature (there are several just in the first chapter of Mackie<sup>22</sup>), and it is very hard to see how there could exist a deductive argument that demonstrates that every possible argument for moral skepticism contains a hidden premise of a particular kind (i.e., a substantive first-order moral claim).<sup>23</sup> I suggest that the lesson to take away from this is that every argument for moral skepticism needs to be grappled with on its own terms. We must always remain open to confronting individual skeptical foxes. Perhaps some such arguments smuggle in, or require the addition of substantive moral premises, but we should not assume this from the get go.

In any case, it is far from clear that Hume's principle is well-suited to provide an explanation of the way in which substantive moral judgments stand in the way of us accepting moral skepticism. Rather ironically, given the origins of Hume's principle, as well as Dworkin's own strategy for employing it, Dworkin's failure to show that Hume's principle is well-suited for this purpose may be best explained by the fact that there is a hidden premise in Dworkin's own reasoning. Shafer-Landau locates this hidden premise: "If nonevaluative claims, by themselves, are unable to vindicate any moral claim [Hume's principle], then they are unable, by themselves, to undermine any moral claim."<sup>24</sup> I agree with Shafer-Landau that this hidden premise is both vital to those of Dworkin's arguments that rely on Hume's principle, and extremely dubious. Here is an example of how an ordinary factual claim can completely undermine a normative claim. Suppose a large meteorite is heading towards earth. Many people think it makes little sense to say *you* ought to stop it (assuming you have no superhero powers, or access to a giant laser, etc.). If, as many people think, the "ought implies can" principle is true – and this is clearly not a substantive moral claim itself – then factual claims surely can, and often do, completely undermine particular moral claims. Although determining the exact modal status of the "can" involved here has

---

<sup>21</sup> DAVID HUME, A TREATISE OF HUMAN NATURE bk. 3, pt. 1, § 1, at 469 (L.A. Selby-Bigge ed., Oxford Univ. Press 2d ed. 1978) (1740). Of course, Hume's target is different than Dworkin's, since Hume is concerned with first-order moralists who suddenly leap to normative conclusions from purely descriptive claims about the nature of God or human affairs.

<sup>22</sup> MACKIE, *supra* note 16, at 15-49.

<sup>23</sup> Regarding Hume's target, one might think one is able to rely simply on induction. However, one cannot rely on induction when it comes to Dworkin's target, because one is hard-pressed to find a single example of a standard argument for skepticism in Dworkin's text where he demonstrates convincingly that the argument depends on a hidden premise of a substantive moral kind, let alone a series of such arguments.

<sup>24</sup> Shafer-Landau, *supra* note 11, at 484.

---

proved to be an interesting task for philosophers, it ordinarily seems quite fine to reason: (1) I cannot  $\phi$ ; (2) if I cannot  $\phi$  then it is not the case that I ought to  $\phi$ ; (3) therefore, it is not the case that I ought to  $\phi$ .

Dworkin seems to be guided throughout Part I of *Justice for Hedgehogs*, and particularly the sections where he appeals to Hume's principle, by an intuition that there is an unbridgeable barrier between our substantive moral judgments and the skeptic's claims. He views this barrier as one that will arise over and over again, whenever a moral skeptic attempts to argue against realism, for the skeptic will always be relying upon inescapable substantive moral judgments, at least implicitly or unconsciously. I share Dworkin's guiding intuition that there is a barrier in the general vicinity of substantive moral judgments, but I think he locates the barrier in the wrong place. He attempts to place the barrier with the moral skeptic – with the skeptic's arguments, and the skeptic's unarticulated premises – when it might be a much better idea to think of the barrier as being squarely on the ordinary moral realist's side. In particular, we can each of us ask: what type of evidence would need to be provided for me to reasonably justify suspending any of my fundamental beliefs about what there is moral reason for me to do (where my belief that there is a reason for me to do a particular action involves me taking it as true that there is a reason to do that action)? I have moved here from speaking of ought claims to speaking of reason claims, because this suits my purposes, but this should not matter so far as enquiry into the objectivity of the normative realm is at issue. The term "reason" can be used in both a normative and a non-normative (merely causal) way, but I wish to restrict my attention to normative usage of the term.<sup>25</sup>

Let me now introduce a substantive thesis about normative reasons, the "reasons as evidence" thesis.<sup>26</sup> This thesis offers an explanation of what a normative reason is. In essence, the thesis states that any claim that some fact is a normative reason for one to do a particular act is, extensionally equivalent to a claim that the same fact is evidence that one ought to do this act.<sup>27</sup> The essence of the thesis can be stated in one sentence:

---

<sup>25</sup> An example of a non-normative usage is "the reason it rained was because the clouds were heavy"; compare this with "she is in pain, so there is a reason for me to help her, even though there is also a reason for me to get home, for a friend is waiting for me there," or "she is in pain, so I have a reason to help her, even though I also have a reason to get home, for a friend is waiting for me there." I will ignore the distinction between "there is a reason for me to . . ." and "I have a reason to . . .," but this distinction is discussed in two papers that I have coauthored with Stephen Kearns. See Stephen Kearns & Daniel Star, *Reasons as Evidence*, in 4 OXFORD STUDIES IN METAETHICS 215, 215-42 (Russ Shafer-Landau ed., 2009) [hereinafter *Reasons as Evidence*]; Stephen Kearns & Daniel Star, *Reasons: Explanations or Evidence?*, 119 ETHICS 31, 31-56 (2008) [hereinafter *Reasons: Explanations or Evidence?*].

<sup>26</sup> This thesis is described and defended in much more detail in *Reasons as Evidence*, *supra* note 25; *Reasons: Explanations or Evidence?*, *supra* note 25.

<sup>27</sup> The thesis also applies to reasons for belief, but I will not discuss this aspect of the thesis here, except to say that it is an advantage of this thesis that it provides a *unified* and

“Reasons as evidence”: Necessarily, a fact *F* is a *reason* for an agent *A* to  $\phi$  if and only if *F* is evidence that *A* ought to  $\phi$ .

Suppose a stranger trips and cries out in pain in front of you. The fact that the stranger cries out is evidence that you ought to help the stranger (does anyone really wish to deny this?). It is also a reason to help the stranger. The “reasons as evidence” account of normative reasons simply asserts that such claims always rise and fall together. A number of objections to this thesis have been made, quite a few of which Stephen Kearns and I have addressed in print,<sup>28</sup> and more have been provided to us since by John Broome<sup>29</sup> and John Brunero,<sup>30</sup> both of whom we expect to respond to in print very soon, but *none* of the criticisms that I have seen so far have anything to do with moral skepticism.

How does the “reasons as evidence” thesis make life difficult for the moral skeptic? As I see it, the skeptic has three choices. The skeptic can: (1) deny that the fact that the stranger cries out in pain is really evidence that one ought to help him; (2) admit that the fact that the stranger cries out in pain is evidence that one ought to help him, but deny that it is a normative reason for one to help him; (3) admit that the fact that the stranger cries out in pain is a reason to help him, since it is evidence that one ought to help him, but deny that it is a *moral* reason. The skeptic should not take option (1). No matter how convinced she is that there is nothing that one ought to do, she should not think that there is not even any evidence that I ought to help someone, or get to work, or eat a meal, etc. This flies in the face of experience. People constantly ask themselves what they ought to do next, and constantly weigh evidence as to what they ought to do. The skeptic could go with option (2), but that is just to deny a part of the “reasons as evidence” thesis, so she will need to carefully consider and critically engage with that thesis (it is beyond the scope of this paper to consider how she might do that).

If she accepts option (3), the skeptic has made it clear that she has a weaker thesis than one might have thought she was defending. The skeptic who takes this approach admits that there are normative facts – facts about what reasons we each have to act in one way or another – and she accepts that these reasons are grounded in evidence concerning what each of us ought to do. She presumably thinks all reasons for action are based in self-interest or prudence, and never in moral demands. Dworkin certainly does not discuss this kind of skepticism, since he does not separate out the objectively normative and the moral. In any case, I believe this type of skepticism is also not compatible

---

*informative* account of reasons for action and reasons for belief, where other accounts fail to do so.

<sup>28</sup> *Reasons as Evidence*, *supra* note 25; *Reasons: Explanations or Evidence?*, *supra* note 25.

<sup>29</sup> See John Broome, Response, *Reply to Southwood, Kearns and Star, and Cullity*, 119 ETHICS 96, 100-03 (2008).

<sup>30</sup> See John Brunero, *Reasons and Evidence One Ought*, 119 ETHICS 538, 538-45 (2009).



with the “reasons as evidence” thesis, in combination with our actual experiences, since we clearly do very often find ourselves confronted with evidence that one ought to help another person for no gain, and often considerable cost, to oneself. Moral demands may not be restricted to such acts, but it is a good place to look for them in our experience. The “reasons as evidence” thesis only needs evidence of particular morally colored oughts to exist in order for it to follow that there are moral normative facts – facts about our moral reasons – and does not need any stronger claim than this.

We can understand the skeptic as being in the business of providing evidence that it is *not* the case that you ought to  $\phi$  (for all and any  $\phi$ ). But you clearly *do* have strong evidence that you ought to help the stranger who is in pain in front of you – the stranger is in pain! You can ask yourself: how strong would the skeptic’s evidence need to be (that is, what credence do you have in the conjunction of her premises, assuming it is perspicuous that she is using a deductively valid argument form, which it may not be) in order to overturn your ordinary evidence that you ought to help the stranger? I would suggest that on any occasion that you ask yourself a question like this one, you will find yourself confronting a huge barrier to accepting the skeptic’s conclusion. Importantly, you can accept the claims of *this* particular paragraph without even accepting the “reasons as evidence” thesis. However, that thesis bolsters the claims of this paragraph, by going further in locating normativity in evidence that one ought to act in one way or another, as well as in oughts.

Notice that evidence that it is *not* the case that you ought to  $\phi$  does not directly correspond to any reasons for action according to the “reasons as evidence” thesis. The right hand side of the essential claim of this thesis, as provided above, contains no room for a “not” before the “ought” – this should not be confused with the fact that it does allow for a “not” after the “ought,” which leaves space for reasons to not  $\phi$ . Evidence that you ought *not*  $\phi$  should not be confused with evidence that it is *not* the case that you ought to  $\phi$ , to echo a point made near the beginning of the paper: a claim that you ought not  $\phi$  is necessarily a first-order substantive normative claim, while a claim that it is not the case that you ought to  $\phi$  is the only one of these two claims that can play a second-order non-substantive role. Hence, the skeptic only ever provides us with *indirect* attacks on our normative reasons, since she is not in the business of providing us with either evidence that I ought to  $\phi$ , or evidence that I ought not  $\phi$ , i.e. evidence that could directly come into conflict with my normative reasons.

In providing us with evidential support for evaluative claims, often of a strong kind, certain ordinary facts, such as that someone in front of me is in pain, provide us with normative reasons, *reasons that are such that the skeptic cannot directly attack them*.

Let me now begin to bring this essay to a close by returning to Hume’s principle, which states that it is illegitimate to infer an *ought* claim from an *is* claim. Recall that Dworkin is wedded to this principle, and although the reader may agree with me that he does not successfully wield the principle against the

---

---

skeptic, he or she may still think Dworkin is right to stand by this principle, and he or she may worry that this principle conflicts with my alternative account of the barrier between substantive moral judgments and moral skepticism.

The most important thing to note here is that Hume's principle was never meant to run contrary to *non-deductive* evidential support for moral claims, and that if it were reformulated to do so it would become much less plausible. Russ Shafer-Landau expresses this point very well when he writes:

Hume never denied that factual claims could nondeductively support or cast doubt on evaluative ones. And that was wisdom on his part, since factual claims can indeed play such roles. Suppose I tell you that a certain killer had no remorse, that his victim was a small child who pled for his life, that the killer was motivated solely by a desire to see his victim suffer . . . . Though this combination of facts does not entail the wrongness of the killing, it certainly offers a great deal of support for such a verdict.<sup>31</sup>

The "reasons as evidence" thesis depends on there being non-deductive evidential support for ought claims, but we should all accept that there is such support, because we are all already familiar with it in our experience: whenever we ask ourselves what we ought to do, we weigh the evidence for acting one way or another, and, to the extent we are deliberating rationally, we do our best to move in whichever direction the evidence most favors.<sup>32</sup>

Dworkin might well respond that there is a deductive argument that involves a normative principle always tacitly lurking behind every judgment of evidential support for a normative claim (i.e., in the present context, behind every judgment concerning a normative reason). Consider, for example, the following candidate for such an argument: (1) A person near me is in pain (non-normative fact); (2) if someone near me is in pain then there is a reason for me to help that person (putative normative principle); (3) therefore, there is a reason for me to help the person near me. If Dworkin were to construe all judgments about our reasons for action in a way that follows this model, he would be using a very broad version of Hume's principle, indeed, and one that

---

<sup>31</sup> Shafer-Landau, *supra* note 11, at 486.

<sup>32</sup> Sometimes we rationally move in one direction or another by responding to evidence concerning what we ought to do, but without thinking explicitly about evidence. This should not be surprising, since, putting cases involving the practical "ought" to one side, we spend much of our lives responding spontaneously to evidence when we form ordinary beliefs, e.g., through perception. I sometimes receive the objection that we do not make many judgments about what we ought to do in our daily lives at all, spontaneously or otherwise. I think this objection relies on confusing the general all-things-considered "ought" with the "ought" or, better, "obligation" or "duty" of moral requirements. Many of our actual ought judgments are very mundane (e.g., the judgment I make when I answer the question "Ought I eat the apple or the banana?"), and we may not always explicitly use the word "ought" (or "should," or some other roughly synonymous term).

Hume himself did not use. In any case, I do not believe that this is the right way to understand my relation to the reason (or evidence that I ought) to help the person near me. I can appreciate and confidently endorse that reason (or evidence that I ought) much more simply and directly than I can appreciate and confidently endorse a normative bridging principle such as (2).<sup>33</sup> This suggests my reasoning does not go via an inference from a normative principle. Furthermore, principles such as (2) do not capture all the normative information that we encounter in our deliberations: in particular, we lose information about the *weight* of various reasons in the above three step argument, and it is very difficult to see how such information could be part of the propositional content of a belief (i.e., built into every bridging principle), or alternatively, how useful it could be to our ethical reasoning to be using bridging principles without such information built into them.

In any case, assuming that we take this objection to be a response that Dworkin might make to Shafer-Landau regarding non-deductive support for normative judgments, we will need to recast the argument that we are supposing lies behind ordinary reasoning, in a way that makes the imagined response by Dworkin much less plausible straight away: (1a) a person near me is in pain (non-normative fact); (2a) if someone near me is in pain then there is evidence that I ought to help that person (putative normative principle); (3a) therefore, there is evidence that I ought to help the person near me. It seems absurd to think that we normally *identify* evidence through using principles like (2a), for we typically *use* evidence in our reasoning, rather than *identify* it by reasoning. Suppose Shafer-Landau decided not to endorse “reasons as evidence”; he could say that (2a) simply does not look like a normative principle at all – it only begins to look like one when it is viewed as equivalent to a normative reasons statement. Suppose instead that Shafer-Landau decided to endorse “reasons as evidence.” He could say that (1a) to (3a) are equivalent to (1) to (3) and that if we initially judge that (1a) to (3a) is a much less plausible example of a type of inference that might be imputed to ordinary agents than the type of inference to be found in (1) to (3), then this judgment about (1a) to (3a) will undermine the imagined objection’s appeal to a broad

---

<sup>33</sup> One reason for thinking this is so is that (2) is not an obviously true principle. To the extent we think of it as a principle, I take it that we very naturally think of it being prefaced by a universal quantifier (i.e., in *all* cases where there is a person standing in front of me in pain, there is a reason for me to help that person), but then we may think there are situations where the relevant reason/evidence is defeated by the presence of other reasons/evidence, if not with the example of pain, at least with a great many examples of reasons that we would ordinarily cite. See generally JONATHAN DANCY, *ETHICS WITHOUT PRINCIPLES* (2004). Alternatively, the principle might contain a *ceteris paribus* clause, but then the argument from (1) to (3) is not strictly valid. Another alternative is that the principle might be formulated in a very narrow way to only apply to situations where other reasons do not interfere, but now the requirement that there always be an implicit inference involved in identifying our reasons is starting to look very mysterious indeed – such inferences now seem redundant.

---

---

version of Hume's principle even more than our judgment about (1) to (3) does.

I do not think the imagined objection from Dworkin that we have been considering has *any* merit, but perhaps it would not really matter to me too much if we did discover that there are always implicit substantive normative principles like (2) involved in all judgments about our normative reasons for action. If this turned out to be true, then we would, in effect, be discovering that there are bedrock normative principles that the moral skeptic will find it extremely difficult to dislodge, and we would then believe that such principles are required for any weighing of reasons for action, whatsoever.<sup>34</sup> On the other hand, if, as I think, there are normally no inferences being made on this level, we have located very basic reasons/evidence that the moral skeptic can never hope to directly attack. Either way, it would still be the case that I have identified a very strong barrier to moral skepticism that explains Dworkin's guiding intuition, without needing to attack the moral skeptic on her own fox-hole covered territory.

---

<sup>34</sup> It might be thought that Dworkin's position that the skeptic himself must be using normative principles could return with a vengeance at this point: suppose the acceptance of any normative reasons at all, whether reasons for action or reasons for belief, always requires an inference that runs via a normative bridging principle; then it would turn out that the skeptic's reasoning always involves substantive principles, and in presenting us with an argument that needs to be assessed (i.e., with evidence for her view), she would be requiring us to use such substantive principles while assessing it. However, the idea that our judgments concerning the reasons for belief we find ourselves confronting *always* depend on inference seems *extremely* implausible. The reason this seems extremely implausible is because it is relatively uncontroversial to understand reasons for belief in terms of evidence, so the point made in the penultimate paragraph of the paper applies. If the reader becomes convinced that reasons for action are also evidence, then the reader should also come to see it as implausible that we typically identify reasons for action through inference.