

---

# DO NORMATIVE FACTS NEED TO EXPLAIN?<sup>1</sup>

BY

JEREMY RANDEL KOONS

---

**Abstract:** Much moral skepticism stems from the charge that moral facts do not figure in causal explanations. However, philosophers committed to normative epistemological discourse (by which I mean our practice of evaluating beliefs as justified or unjustified, and so forth) are in no position to demand that normative facts serve such a role, since epistemic facts are causally impotent as well. I argue instead that pragmatic reasons can justify our continued participation in practices which, like morality and epistemology, do not serve the function of causal explanation.

The term ‘epistemology’ refers to two rather different practices. Construed narrowly, it is the attempt to construct a set of formal conditions for the justification of beliefs (e.g., explanatory coherence among beliefs; proper inferential relation to foundational beliefs; proper causal genesis, à la reliabilism, etc.). Construed more broadly, it is the practice, common in every field from chemistry to philosophy, of evaluating beliefs as justified or unjustified, scientific methods as rational or irrational, etc. Philosophers such as Rorty have argued that we should quit practicing epistemology in the narrow sense, but even philosophers convinced by Rorty could agree that epistemic discourse in the broad sense is an essential aspect of every branch of knowledge.

Although most philosophers are convinced of the viability of epistemology in this broad sense, many are somewhat less sanguine about the prospects for normative *moral* discourse. Some of this skepticism stems from an influential objection raised by Gilbert Harman in the first two chapters of his 1977 book *The Nature of Morality*. According to Harman, moral facts neither figure in the best causal explanation for any observations, nor are reducible to predicates which are explanatory. Harman infers from this that there are no moral facts, and that some form of moral nihilism must therefore be true.

While Harman's explanatory requirement (as it has been named) has caused a good deal of worry concerning moral discourse, fewer philosophers have worried about its ramifications for epistemological discourse. The goal of this paper is to use the average philosopher's commitment to epistemological discourse as a tool to rescue moral discourse from this particular objection. In other words, I wish to argue that philosophers who use epistemological discourse in the broad sense – and this is the sense with which I will be concerned in this paper – are in no position to demand that moral facts serve a causal role in explaining observations.<sup>2</sup> I would like to suggest, rather, that those who are committed to epistemological discourse must concede that pragmatic reasons can serve as a bulwark against eliminativism, even if the discourse in question does not provide genuine causal explanations, nor reduces to causally efficacious facts. Finally, since Harman thinks that epistemology might be saved by reducing it to natural facts, I will argue that there is no in-principle barrier to the reduction of moral facts which does not also serve as a similar barrier to the reduction of epistemological facts. Essentially, then, what I am presenting here is an equality argument: I am arguing that epistemology and morality stand and fall together. Each is as robust (or feeble) as the other. Thus, those philosophers – such as Harman, Mackie and Ayer – who believe in a robust epistemology but who are eliminativists or nihilists about morality have it badly wrong.<sup>3</sup>

I hope, too, that this paper helps to clear up some other worries philosophers have about the status of moral discourse. For eliminativism and full-blown moral realism are not the only possible stances to hold regarding moral discourse. Far more interesting and important are the intermediate positions vis-à-vis moral discourse held by legions of philosophers. These are the views of philosophers such as Blackburn, Gibbard, Wright, and others, who endorse a sort of “quasi-realism.” According to them, moral utterances can be true, and the true ones even have some form of authority over us; but these moral utterances aren't as robustly true as, say, true scientific utterances. But there is a residual unease in adopting this middle position. “Quasi-realism” (if I may hijack Blackburn's term) brings with it the worry that moral discourse is second best; it may still be alive, but it only limps along. But the conclusion of this paper should alleviate some of these concerns. After all, if I am right, then moral discourse is on par with something as robust as epistemological discourse. Epistemological discourse is the very thing philosophers spend their life employing; it is their stock in trade. After all, what philosophers do is advance *and evaluate* arguments; philosophers try to tell us what it is *rational* to believe about ethics, about the meaning and reference of words, about mental states. Without epistemology, philosophy as we know it grinds to a halt. Science does the same: how can science succeed if there is no distinction between science and pseudoscience, rational and irrational

research methods, justified and unjustified scientific beliefs, and so on? I will have much more to say on this topic later in the paper. Let us begin.

### *The explanatory requirement*

Harman begins his criticism of moral realism by asking the question, "Can moral principles, like scientific ones, be tested against the world?" The first response given to this question is that we do, indeed, observe moral properties, in the sense that we can have a perceptual experience (say, hoodlums lighting a cat on fire) and form (non-inferentially) the judgment, "That is cruel." Harman writes

If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can *see* that it is wrong. (Harman 1977)

But do we perceive moral properties in the same way we perceive scientific properties and objects? Many point out that we perceive moral properties only because we have been trained to do so. But this holds of scientific properties, as well: it is fairly well accepted in philosophy these days that there is no such thing as an uninterpreted observation language.<sup>4</sup> There is no theory-neutral observation language, and theories are the sorts of things that are taught and learned. However, Harman thinks he has another tool for prying apart moral and scientific perception: the explanatory requirement.

It is a current tenet of epistemology that we are justified in positing the existence of an entity if this entity is invoked in the best explanation of some accepted phenomenon. For example, say a scientist sees a vapor trail in a cloud chamber. The best explanation for this observation is that a proton just went by; therefore, we get to be realists about protons.

Now let us consider the case of the cat. What do we need to invoke to explain our observation that igniting the cat was cruel? Certain psychological facts, facts about upbringing, etc., enter into the equation, but, Harman argues, cruelty need not. The best explanation of the observation does not invoke cruelty or posit moral properties; thus we don't get to be realists about morality.

Clearly, Harman is on to something. I am willing to concede that moral facts can never causally explain non-moral facts. This is actually a matter of dispute; some argue the contrary (see, for example, Brink 1989, pp. 182–97). I don't, however, wish to enter into this debate. My case is hardly hurt if the naturalists are right: the more morality resembles scientific discourse, the less plausible is an eliminativist argument. So instead, I wish to take the weaker position and argue that even if moral facts

don't figure in the best explanation of any observations, we *still* need not be eliminativists about morality. If causal explanatory impotence is a sufficient reason to jettison morality, then much else ends up going with it. I will argue that if we take Harman's requirement seriously, we will have to jettison much discourse we should be reluctant to discard.

#### MORALITY AND CAUSAL EXPLANATION

To be an explanatory eliminativist about a particular set of entities or sentences is to deny that those entities or sentences are genuinely explanatory.<sup>5</sup> Do moral facts enter into causal explanations? Harman suggests not, and without giving an argument, I am willing to concede that they do not. While we might casually employ moral facts in workaday explanations ("The overseer's cruelty ultimately led to the slaves' revolt"), these moral facts are eliminated in a more penetrating analysis ("The overseer's treatment of the slaves caused them great suffering, which in turn bred rampant resentment among the slaves, leading to their revolt"). Moral facts certainly do not appear in the most penetrating causal explanations (i.e., scientific explanations).

There is an important misunderstanding which might arise at this point. Many will claim that moral facts do enter into some causal explanations: they enter into rational explanations. I concede this, but will argue that moral facts enter into rational explanations in a way that is consistent with the causal impotence of these facts.

It is true that rational explanations are causal. Some would deny that they are, but I must confess some puzzlement regarding this denial. After all, if our beliefs and desires don't cause our actions, then in what sense are we in control of our external behavior? When I give a rational explanation – "He went to the refrigerator because he wanted a glass of milk, and believed that there was milk there" – the beliefs and motivating states we posit in this explanation must be causally efficacious, or else we haven't provided an *explanation* of his action (which, whatever else it may be, is a physical phenomenon) at all. If we deny that beliefs and desires may cause action, we are forced into some strange Malebranchian occasionalism or Leibnizian parallelism if we are to explain the correlation between his desire for milk and his action (going to the refrigerator). I shall therefore assume that rational explanations are causal ones.

I assume the reader has an intuitive idea of the shape of rational explanation: these explanations causally explain agents' actions in terms of reasons. Moral facts figure in some rational explanations. Thus, while not all rational explanations involve moral properties, some such explanations will: "I returned the lost wallet because it was the right thing to do." It would be misleading, though, to think that the action's rightness caused me to return the wallet. Harman, for one, would deny that

this is the case. Furthermore, I am inclined to agree with him. I returned the wallet not because it was the right thing to do, but because *I believed* that it was the right thing to do. Thus, Sellars writes that “obligation enters into the causal order only as an element in the intentional object of a mental act,” that is, “*via* facts of the form *Jones thinks (feels) that he ought to pay his debt*” (Sellars 1953, p. 222). Thus, moral properties and facts enter only into rational explanations, and only as the objects of intentional states.

I think anyone will concede that this degree of causal involvement is not sufficient to vindicate moral realism; *anything* can serve as the object of an intentional state. After all, it might well be true that Jones *believes* that his chronic headaches are caused by the tiny unicorns which dance on his forehead every night while he sleeps, and this belief may even cause Jones to take certain actions (e.g., grease his head so the unicorns lose their footing), so the belief is causally efficacious. Even if this belief-attribution is true, we shall still be eliminativists about tiny dancing unicorns (I hope).

Of course, not all explanations are causal. Consider two examples from Lance and O’Leary-Hawthorne (1997, p. 386): there are infinitely many primes because for any number  $n$ , the result of adding 1 to the product of the first  $n$  primes is also a prime; and “We do not know who will win the election because given the sample size of our study, its margin of error was greater than the predicted vote differential.” Both of these seem to be non-causal and also genuinely explanatory – they not only give reason to believe the claim in question, but show why the claim is true. And as van Fraassen, for example, has convincingly argued, explanations are those propositions that answer “why” questions.

An important category of non-causal explanation is the category of normative explanations. When we give a normative explanation, we explain how an action or feature alters or contributes to a particular normative characterization of an action, person, or object. For example, while watching a football game I might say, “His crossing the 35-yard line while carrying the ball made it first down.” The football player’s crossing the 35-yard line does not cause the first down; rather, it changes the normative status of the players in the game. True, his crossing the 35-yard line causes certain actions which accompany the awarding of a first down – movement of the down markers, etc. – but the arrival of a first down is itself a normative occurrence, not a causal one. In the moral case, I might say, “That was cruel because it caused her pain.” Her pain does not cause cruelty; nor does the cause of her pain cause cruelty. The “because” in this explanation is a normative, not a causal one. The same is true in epistemology, even the versions that rely on causal chains: if I say that belief B is justified because it was caused in the right way, I am asserting that the belief, not the justification, was caused. Rather, the

justification is an additional, normative fact which is not *caused* by the belief, or by the cause of the belief. If this sort of epistemological theory is right, a belief's justification is explained by that belief's causal etiology, but this is a normative sort of explanation, not a causal one. Examples of normative (and hence non-causal) explanation can be found in any normative endeavor, from epistemology to morality.<sup>6</sup>

To sum up: we have seen there are causal and non-causal explanations. Moral facts figure in certain non-causal explanations, to wit, in some normative explanations. Morality enters the causal order through rational explanations, but only as the object of an intentional act or state. Rational explanations in which morality seems to play a more robust causal role (e.g., "I returned the wallet because it was the right thing to do") are more perspicuously rendered (e.g., "I returned the wallet because I believed it was the right thing to do") to reflect the causal impotence of moral facts.

Thus, I am willing to concede the premise of Harman's objection: morality does not enter, in any interesting way, into causal explanations. That is to say, I concede causal explanatory eliminativism with regard to moral discourse. Does this imply *doctrinal eliminativism* about morality, i.e., the claim that all moral utterances are false, or that they are neither true nor false? If it does, then epistemology goes with it, because this causal impotence is echoed by normative epistemological discourse.<sup>7</sup> That is to say, epistemological properties are equally impotent. I argued above that beliefs can be causes. However, justification and warrant are not themselves causally efficacious. All the same moves which can be made to deny that moral properties are causally efficacious can be made to deny that epistemological properties are causally efficacious. So we can say, "I believe P because P is justified," but just as in the moral case, the real causal explanation is: I believe P because I believe that P is justified. The same move can be made for prudential explanation, too ("I did Y because it was optimal" becomes "I did Y because I believed it optimal"). So if causal explanatory impotence entails doctrinal eliminativism, then epistemology and rationality are eliminated as well.

Epistemic facts are causally impotent even if epistemology is naturalized, as writers such as Alvin I. Goldman have attempted to do. For these philosophers, justification of knowledge consists in the appropriate sort of causal connection to natural facts. But notice that it is the natural facts which are causally efficacious; the epistemic facts are themselves causally impotent. The epistemic facts consist of certain causal relations, but causal relations don't themselves cause anything; it is the *relata*, not the relation, which are causally efficacious. For example, suppose on a pool table, the cue ball strikes the eight ball and causes it to roll into the corner pocket. Thus, a causal relation obtains between the two pool balls. Notice, though, that the causal relation is not itself causally efficacious.

If it were, I suppose perception of causal relations would be unproblematic, and Hume would never have become a skeptic about them. Rather, the relata (the cue ball and the eight ball) can be causes; the relation cannot. So epistemic facts are causally impotent even if epistemology is naturalized.<sup>8</sup>

Clearly, we may colloquially employ phrases in which moral and epistemological facts appear to serve a causal explanatory role. Why did Selim Woodworth fail to rescue the Donner party? Because he was no damned good.<sup>9</sup> Why did you use a double-blind study? Because that is the best way to determine drug toxicity. Harman's point (although he only makes it about morality) is that the moral/epistemological part of the explanation is eliminable, and does not appear in a more penetrating natural-scientific explanation.

We can thus see that the cost of adopting the explanatory requirement would be high: we would be deprived of epistemological discourse as well as moral discourse. We will see in a moment whether such pragmatic considerations are relevant in the realism-antirealism debate.

### *Eliminativism and pragmatic reasons*

We have so far concluded that causal explanatory eliminativism obtains with regard to morality and epistemology. In other words, neither moral nor epistemological facts are genuinely causally explanatory. What does this show about morality and epistemology? I think it shows at most that *ontological eliminativism* obtains with regard to these types of discourse (although I will neither insist on this conclusion nor argue for it; I am willing to concede it for argument's sake). To be an ontological eliminativist about a particular subject matter is to claim that certain noun-like words or phrases belonging to that subject matter do not refer. Most of us are ontological eliminativists about unicorns and phlogiston; many are ontological eliminativists about holes, shadows, and the like. Thus, I am willing to concede something like the following entailment: if K-talk is not involved in the best causal explanation of any uncontroversial phenomenon, then Ks do not exist. The mistake is thinking that failure to satisfy the explanatory requirement entails doctrinal eliminativism. Recall, from the introduction, that doctrinal eliminativism holds that none of the sentences belonging to a particular subject matter are ever true, either because they are all false, or neither true nor false. Consider the sentence, "The average American mother has 2.3 children." It is uncontroversial, I take it, that there is no entity to which this sentence refers; there is no such thing as the average American mother.<sup>10</sup> Indeed, I imagine that there are no causal explanations in the language of austere physics that make appeal to this entity. So we have an example of a

sentence which has the following two properties: the sentence uses a non-referring noun-like expression (which expression fails to satisfy the explanatory requirement), and the sentence is uncontroversially true. Hence, causal explanatory eliminativism does not entail doctrinal eliminativism. In other words, K-talk can be true even if there are no Ks.<sup>11</sup> (Examples can be multiplied indefinitely. Consider “John has a big ego.” We may concede that this phrase conveys information, and can be true, even if we are ontological and causal explanatory eliminativists about egos.)<sup>12</sup>

#### EPISTEMOLOGY AND PRAGMATIC REASONS

What, then, is the correct test for doctrinal eliminativism? Perhaps we can discover this by asking what sort of doctrinal eliminativism test epistemology can survive. I want to articulate the right test by showing what test epistemic discourse in the broad sense passes. I think this is an appropriate strategy, since we know we want to keep epistemology in this broad sense. Having articulated the standard which is met by our normative epistemic discourse, I want then to suggest that morality meets this standard as well. Epistemology, like morality, is not an explanatory project: epistemological properties (justified, rational, etc.) are not causally efficacious. What sorts of considerations support our epistemic practice, then? In virtue of what, precisely, are epistemic claims doctrinally ineliminable? I want to argue that they are ineliminable because they are an essential part of a larger, successful *pragmatically* justified project.

Let me start by discussing Thomas Kuhn’s article “The Essential Tension” (1959). In this article, he notes that we typically think that one of the chief virtues of a successful scientist is the ability to engage in so-called “divergent thinking”: the ability to look at obvious facts without necessarily accepting them; the ability actively to imagine unique possibilities, instead of merely interpreting facts in the way taught in textbooks and by scientific authorities. Kuhn acknowledges the importance of divergent thinking; without it, there would be no scientific revolutions. Ptolemaic astronomy would never have been discarded in favor of Copernicanism; oxygen would never have dislodged phlogiston; and Einstein would not have supplanted Newton. But very little of science is revolutionary; “normal” science is more akin to “puzzle solving” than to revolution.<sup>13</sup> And in normal science, Kuhn emphasizes, “convergent” thinking is more useful than its divergent counterpart.

[N]ormal research, even the best of it, is a highly convergent activity based firmly upon a settled consensus acquired from scientific education and reinforced by subsequent life in the profession. Typically, to be sure, this convergent or consensus-bound research ultimately results in revolution. Then, traditional techniques and beliefs are abandoned and replaced by new ones. But revolutionary shifts of a scientific tradition are relatively rare, and extended periods of convergent research are the necessary preliminary to them. (227)



The textbooks of an era indoctrinate scientists-in-training with the current paradigms. Kuhn notes that in the latter half of the 19th century, textbooks stated unequivocally that light was wave motion. Before that, books on optics were equally unequivocal in their assertions that light consisted of particles. And of course now textbooks state (again, unequivocally) that light exhibits properties of both waves and particles. It may seem odd that textbooks and scientists are so dogmatic in asserting theories which ultimately prove false, but Kuhn argues that science would not have achieved its current lofty stature were it not for its emphasis on convergent thinking.

The history of theories of light does not . . . begin with Newton . . . From remote antiquity until the end of the seventeenth century there was no single set of paradigms for the study of physical optics. Instead, many men advanced a large number of different views about the nature of light . . . As a result, a new man entering the field was inevitably exposed to a variety of conflicting viewpoints; he was forced to examine the evidence for each, and there was always good evidence . . . This earlier mode of education was obviously more suited to produce a scientist without prejudice, alert to novel phenomena, and flexible in his approach to his field. On the other hand, one can scarcely escape the impression that, during the period characterized by this more liberal educational practice, physical optics made very little progress. (231)

This preparadigm or preconensus period in optics is echoed in virtually every other discipline, from dynamics to physiology. And further, the study of the preparadigm periods of these other disciplines strongly suggests that “without a firm consensus, this more flexible practice will not produce the pattern of rapid consequential scientific advance to which recent centuries have accustomed us” (232). Without a theory to which one is committed, one does not know what phenomena are significant, and what problems are worth solving (or perhaps indeed what the problems are in the first place). Further, one must be committed to a theory to undertake the serious work which is required to extend, deepen (*and ultimately overthrow*) the theory. “Who, for example, would have developed the elaborate mathematical techniques required for the study of the effects of interplanetary attractions upon basic Keplerian orbits if he had not assumed that Newtonian dynamics, applied to the planets then known, would explain the last details of astronomical observation? But without this assurance, how would Neptune have been discovered and the list of planets changed” (235)?

How does any of this bear on the project of epistemology? The answer is this: epistemic discourse is one of the tools that enables a consensus to arise and be maintained, and is therefore one of the tools which speeds scientific progress, which is something we value. *The job of epistemology just is to sort beliefs into those that are to be believed, and those that are not.* For a physicist working in 1800, it was rational to believe that

Newtonian physics was correct. It was rational to carry out research projects using Newtonian techniques, and it would have been irrational to attempt to use Aristotelian physics to account for planetary motion. A physicist who incorporated unjustified techniques (such as those of Aristotelian physics) would have been sanctioned by partial or complete exclusion from the profession of physics. Those who were adept at applying Newtonian techniques to the study of planetary motion were rewarded by inclusion.

Epistemic evaluation can be discursive or practical. In other words, we can *call* a methodology irrational or a belief unjustified, or we can *carry out* our scientific practice in a way which favors this methodology, or which presupposes the falsity of that belief. And when our 18th-century professional physicists applied Newtonian techniques to the study of planetary motion, they were practically (though not discursively) endorsing these techniques. Similarly, the exclusion of the Aristotelian physicists from their ranks was a form of practical evaluation; it was an action which practically judged belief in Aristotelian methods to be unjustified. This practical evaluation is crucial to scientific progress. As Kuhn pointed out, if we don't practically favor a particular paradigm over all others, science does not progress rapidly. Rapid scientific progress requires practical epistemic evaluation. Hence, epistemology serves the pragmatic end of hastening scientific progress.

It might be objected that epistemic practice, even in the broad sense, is not the sort of thing that can be implicit in scientific practice. Epistemology in this sense refers to the labeling of beliefs as justified and unjustified, the labeling of methods as rational and irrational, and this labeling must be explicit (spoken, written, etc.). Hence (goes the objection), something other than epistemology does the primary work of ensuring scientific consensus.

But evaluative beliefs, such as those of epistemology, are surely the sort of thing that can be implicit in practice. For example, a person may firmly believe that lying is wrong, even if she never utters the sentence, "Lying is wrong." Her belief can be expressed by her constant truth-telling, especially in those cases where lying would clearly be more expedient for her. Similarly, my belief that the dishwasher will wash my dishes is implicit in my actions (placing dirty dishes in the washer and turning it on) even if I don't utter the sentence, "My dishwasher will wash my dishes" and don't use the belief in a bit of explicit internal practical reasoning. We can evaluate or believe on particular occasions without saying anything, even *in foro interno*. Our evaluation and belief is internal in some cases, and implicit in others. In fact, I imagine that the majority of our practice occurs at the level of the implicit – we only proceed to the level of explicit discourse when we need to inform, educate, or argue.<sup>14</sup> So it is with epistemological evaluation. Our belief that double-blind studies

are more reliable than chicken bones is implicit in our decision to perform a double-blind study instead of visiting a practitioner of bone-reading. Thus, epistemology is a crucial element of scientific progress: without convergent practical epistemological evaluation implicit in scientific practice, science advances only slowly. We value rapid scientific advancement. Hence, epistemology as a practice is justified pragmatically.<sup>15</sup>

Epistemic discourse serves other uses – all pragmatic. At the level of explicit discourse, it can allow for revision of our scientific and other sorts of truth-seeking practices, thereby serving an auxiliary role to such practices. For example, I might suggest to you that instead of eating peyote in an attempt to contact the spirits and find out from them whether the sun orbits the earth, you might instead read the Bible and find the answer there. “Why should I do that instead of eating peyote?” you inquire. I reply that reading the Bible is more likely to lead you to the truth on this matter, or that the Bible contains true things and peyote just makes you hallucinate. In other words, I say something which has the same practical significance as the epistemological sentence, “Consulting the Bible is a more rational method of investigating scientific matters than is the consumption of peyote.” I am communicating to you, either practically or discursively, my belief that consulting the Bible is a more (epistemically) sound method than is eating hallucinogens.

Consider another use of epistemic evaluation. It allows me to know what sort of practices confer commitments and entitlements on me. For example, if you tell me that Stephen Hawking is a competent physicist, I know that I am entitled to believe his pronouncements on physics. Furthermore, I am under an obligation or commitment to defer to Hawking on questions of physics. Again, this sort of evaluation can be discursive (you can say “You are justified in believing what Stephen Hawking says about physics”), or practical (suppose that, as a matter of practice, you defer to him on questions of physics).<sup>16</sup>

#### MORAL ELIMINATIVISM AND PRAGMATIC REASONS

It appears, then, that our epistemic practice is pragmatically justified. Just as particular scientific beliefs are justified by being part of a larger, successful explanatory project, our individual epistemic judgments are justified in virtue of belonging to a larger, pragmatically-undertaken project. This is an important result, for (assuming we are not yet eliminativists about epistemology) it allows us to appeal to pragmatic considerations in defending a discourse against doctrinal eliminativism (recall that doctrinal eliminativism is the thesis that the sentences of a discourse are all false, or neither true nor false). And of course, in the absence of explanatorily potent epistemic and moral facts, this is exactly

the sort of considerations we must appeal to to justify our epistemic and moral practice.

Certainly, not all pragmatic justifications are created equal, just as not all explanatory justifications are created equal (phlogiston fails; oxygen succeeds). Not all sorts of discourse are worthy of being pursued. The goal “to maintain the political power of the clergy” would not be a particularly good reason to continue pursuing a certain type of discourse. However, the goal “to facilitate human interaction” might be an excellent goal, as is “to explain the natural world.” The point is, a question of whether a type of discourse is robust and worthy of pursuit is not a question which can be decided on *a priori* grounds, merely by gesturing to natural facts. This is a question which is decided on an *a posteriori* basis, on the basis of reasons (among them *pragmatic* reasons).

I should also note that not all *means* to pragmatically justified ends are created equal, either. While evaluating the intentions of a potential suitor is probably worth doing, consulting a psychic is not an effective means to this end. It does not serve the pragmatic end we want it to serve. Moral discourse, however, clearly seems to facilitate human interaction. It is an effective means to the pragmatically justified end of interacting safely and profitably with our fellow humans. Epistemological discourse seems to be an effective tool for advancing our scientific and other truth-seeking practices. Thus, since both morality and epistemology are effective means to worthwhile ends, we should continue to practice both.

I do not think that I can overemphasize the role of pragmatic reasons in questions of doctrinal eliminativism. Not only is it wise to consider pragmatic reasons in these debates, but our actual practice recognizes the importance of such reasons. Consider an example of how pragmatic reasons argue in favor of continuing a particular practice.<sup>17</sup> The rules of games such as basketball are human inventions; no one argues (I hope) that they are answerable to basketball facts in the world. It is important to note, though, that no one argues on these grounds that we ought to quit playing the game of basketball. What would be the point of quitting? (Note that when we ask what the *point* of a move would be, we are asking a *pragmatic* question.) Would we be better off not playing basketball? Of course not. Notice how out of order, and indeed bizarre, we would regard a Harman-style attack on the practice of basketball (basketball rules don't explain observations, aren't reducible to natural facts, etc.). Such objections are beside the point, because explanation of the natural world is not part of the aim of basketball. The point of basketball is to have fun, and perhaps also to get exercise. Basketball achieves the pragmatic ends at which it aims, and since these ends are themselves acceptable ones (remember, not all pragmatic ends are worthy of pursuit), we have good reason to continue the practice of basketball.

The case of basketball is analogous to the moral case: moral facts don't stand in causal, explanatory relation to natural facts, but that's not the point of moral discourse, anyhow. Moral discourse serves an end other than explanation, and we would not be better off eliminating morality. Thus, Harman-style challenges are just as inappropriate to our moral practice as they are to our practice of basketball.

I would argue that morality survives the *a posteriori* challenge. After all, would we be better off not playing the moral game? It seems likely that the answer is "No."<sup>18</sup> Just as the rejection of epistemology would be tantamount to giving up the ability to sort beliefs between those to be believed and those not to be believed, so would the elimination of morality be tantamount to giving up the ability to sort actions between those to be done and those not to be done. Since pragmatic considerations can defend theories against doctrinal elimination, and since moral discourse acquits itself admirably on pragmatic grounds, I should think that we have an excellent reason to be doctrinal realists about moral discourse. Thus, the justification I offer for morality is a sort of "decynicized Hobbesianism."<sup>19</sup> We have interests – selfish *and* other-regarding – and these interests of ours are best met if we enter into a cooperative arrangement, morality.

Given the course of the argument, it might seem odd that I attempt to ground morality and epistemology in what are, at the end of the day, prudential reasons.<sup>20</sup> After all, I have been arguing that moral, epistemic, and prudential reasons are all equally hardy – none is more or less robust than the others. But how can I maintain both that (a) morality and epistemology are grounded in prudential reasons, and that (b) prudential reasons are no hardier or more secure than moral and epistemic reasons? I think we need to distinguish between (1) a discourse's being hardier or more secure than another and (2) a discourse's being logically prior to another. Prudential reasons are logically prior to moral and epistemic reasons. We start out with prudential reasons – they are there from the beginning – and the other sorts of reasons arise out of these prudential reasons, but moral and epistemic reasons are no less hardy than the prudential reasons out of which they arise. Consider an analogy: suppose we want to prove a logical or mathematical theorem, T2. As a preliminary step, we need to prove T1; T1 is an essential element in the proof of T2. Thus, T1 can be said to be logically prior to T2, in the sense that the proof of T2 rests on T1. But we should not infer that T1 is better grounded than T2; both theorems are certain and necessary. So logical priority does not translate into greater justificatory robustness. In any case, the thesis of the paper is that morality and epistemology are on a par; since both rest on prudential reasons, my thesis is not harmed if it turns out that prudential reasons are better-grounded than moral or epistemic ones.

For the reasons just rehearsed, though, I don't think prudential reasons are better-grounded.

Some philosophers will no doubt claim that we ought to continue to use moral talk, but with our fingers crossed, so to speak. That is to say, we should continue to use this discourse while at the same time acknowledging that the sentences of this discourse aren't *really* true. But if this is all that doctrinal eliminativism amounts to, then it amounts to very little. If doctrinal eliminativism about moral discourse concedes that we ought to continue practicing moral discourse, then doctrinal eliminativism is itself a fairly empty doctrine. A philosophical critique of a particular practice isn't particularly damaging if, in the end, it concludes that we ought to carry on pretty much as before. And in any case, morality and epistemology would be equally deflated, so the thesis of my paper – that morality and epistemology are equally objective – remains undamaged.

Harman's objection, thus, is off the mark because the *point* of moral and epistemic discourse isn't to explain (and as I noted above, when we talk about the point of a practice, we are talking in pragmatic terms). Harman's objection might be relevant if directed at a discourse which offered allegedly scientific explanations, but we need not regard this as the pragmatic end served by moral and epistemic discourse. And we ought not adopt Harman's explanatory requirement as a test for doctrinal eliminativism because the pragmatic cost of such a move would be too high. We would have to abandon every practice which served an end other than explanation. Morality, epistemology, and other areas would be eliminated. Would we be better off then? I have argued that we would not be better off if we did not practice epistemology; and I will assume that we are better off practicing morality than not. Thus, we ought to continue practicing morality and epistemology.

There is an important lesson to be gleaned from all of this. A few pages back, I wrote that since Harman's explanatory requirement was unsuitable to serve as a criterion for doctrinal eliminativism, I needed to propose a criterion of my own, so as not to open the floodgates to all sorts of unacceptable discourse (astrology, tea-leaf reading, etc.). It now seems clear to me, however, that no such criterion can be suggested which will apply to all cases. We cannot say *a priori* what the correct connection is between acceptable sorts of discourse and the world. Rather, in each case, scientific *and pragmatic* reasons will apply in varying degrees, with scientific reasons in general serving a subsidiary role to the pragmatic ones. In astrology, scientific reasons might tell us that the *pragmatic* goals of astrology (namely, accurate prediction of the future) are not attainable by study of the stars and planets, but science cannot evaluate the goals of astrology. The same applies in morality: science might tell us whether moral practice serves the goal of pleasant coexistence, but it

cannot evaluate the goal itself. Scientific explanation is itself a project we carry out for pragmatic reasons (thirst for knowledge, instrumental control over the world, etc.).<sup>21</sup> We cannot say in abstraction from individual cases what precise role will be played by scientific and pragmatic reasons, or what the precise mix of the two will be. Thus, questions about doctrinal eliminativism are not ones which can be addressed *a priori* or according to a formal set of rules, but ones which must be answered as part of an ongoing debate concerning the appropriateness, objectivity and usefulness of various types of discourse. It is true that scientific claims generally stand in some causal relation to natural facts, but this causal relation is only explanatory and justificatory in the context of a larger, successful explanatory project, a project whose undertaking is warranted not by scientific reasons, but by pragmatic ones. The fact that a discourse stands in some causal relation to a certain class of natural facts does not in itself explain why that discourse is more worthy of pursuit than other ones which have different goals.

At this point it might be objected that the pragmatic goals of morality and epistemology are not equal. Morality is concerned with peaceful, flourishing coexistence or *somesuch*, but epistemology aims at *truth*. Truth is something different altogether; truth is something that a creature has to pursue to count as an agent at all. This seems to set the practice of epistemology apart from other endeavors; it seems to give epistemology a metaphysically necessary justification. This is no mere pragmatic justification at all (goes the objection). Hence, epistemology and morality are not on a similar footing at all; epistemology has a more secure foundation.

This is a powerful objection, but it rests on a mistake. The critical thing to notice is that epistemology is not aimed merely at truth.<sup>22</sup> After all, if we *only* cared about truth, then we could merely sit by and calculate arithmetic tables all day long, thereby providing ourselves with an unlimited supply of true (indeed, necessarily true) beliefs. What this example shows is that truth is too thin of a goal to guide epistemology all by itself. We care about more than just truth; we want simplicity, explanatory coherence, instrumental control and predictive accuracy, and so forth. This is why I do not think our epistemic practice can be given the metaphysical grounding suggested by the objection; truth is not the only goal of epistemology,<sup>23</sup> and the goals of epistemology are no more worthy of pursuit than the goals of morality.

### *Worries about pragmatic justification*<sup>24</sup>

Probably the most difficult objection to the pragmatic justification of a practice is that it conflates theoretical and practical justification. Theoretical justification justifies the content of a belief, whereas practical

justification justifies the act of holding a belief. The worry is that pragmatic reasons cannot justify the content of a belief (how could they?), and if we allow them to justify beliefs practically, then we run the risk of holding practically justified, yet false, beliefs. Let us turn to a detailed discussion of this worry. First, I will show why pragmatic reasons cannot be used to justify individual beliefs. I will next show the constraints that must be placed on the pragmatic justification of practices.

#### PRAGMATIC JUSTIFICATION AND INDIVIDUAL BELIEFS

Consider an individual belief that, although false, might be justified pragmatically.

[S]uppose I am an Auburn football fan believing that Auburn will win the Alabama game this year. Is this belief justified? Theoretically, no; a comparative analysis of games played by the two teams this season makes this all too clear. My belief is just wishful thinking. But practically, maybe yes. After all, I might need to believe that Auburn will win the Alabama game to be able to maintain any interest in yet another mediocre season, and to get enjoyment from watching the team's other games. So then it is a good thing that I have the belief – and if I did not, I would have a reason to acquire it.<sup>25</sup>

But surely this type of belief-justification is not permitted. Fortunately, I need not concede that it is permitted. When I say that we might pragmatically justify practices, I am not thereby saying that we may appeal to pragmatic considerations when justifying individual beliefs within the practice. Indeed, the reasons that serve to justify the practice itself are not allowed to “infiltrate” the practice and justify individual beliefs within that practice; the practice ceases to function if that is allowed to happen. Simply put, the goals of the practice are not met if you allow the pragmatic justification of the practice to trickle down to the level of individual acts or beliefs. This was seen clearly by Rawls in his article “Two Concepts of a Rule.” In this article, he was trying to establish the superiority of rule utilitarianism over act utilitarianism. With act utilitarianism, utilitarian considerations justify individual actions. Unfortunately, since act utilitarianism is plagued by free rider problems and the like, it in fact fails to maximize utility. This is where rule utilitarianism comes into play: utilitarian considerations justify the adoption of a practice (consisting of a set of moral rules), but individual actions are justified by appeal to the rules of the practice, and not by appeal to considerations of utility. Here is Rawls:

There are obvious utilitarian advantages in having a practice which denies to the promisor, as a defense, any general appeal to the utilitarian principle in accordance with which the practice itself may be justified. There is nothing contradictory, or surprising, in this: utilitarian



(or aesthetic) reasons might properly be given in arguing that the game of chess, or baseball, is satisfactory just as it is, or in arguing that it should be changed in various respects, but a player in a game cannot properly appeal to such considerations as reasons for his making one move rather than another. It is a mistake to think that if the practice is justified on utilitarian grounds then the promisor must have complete liberty to use utilitarian arguments to decide whether or not to keep his promise. The practice forbids this general defense; and it is a purpose of the practice to do this. Therefore what the above arguments [against utilitarianism] presuppose – the idea that if the utilitarian view is accepted then the promisor is bound if, and only if, the application of the utilitarian principle to his own case shows that keeping it is best on the whole – is false. The promisor is bound because he promised: weighing the case on its merits is not open to him. (1955, p. 16)

Thus, in acknowledging that pragmatic considerations can justify practices, we need not concede that pragmatic considerations can justify individual beliefs, any more than we need to concede that the baseball or chess player may appeal to utilitarian considerations to justify a particular play or move, or any more than a utilitarian need concede that a person may appeal directly to considerations of utility to justify a particular action. What this objection fails to recognize, then, is the difference between the pragmatic justification of a practice and the pragmatic justification of an individual belief. The former is permitted; the latter is forbidden.

#### PRAGMATIC JUSTIFICATION AND PRACTICES

The critic is not silenced. He will concede the following: I (the author of this paper) am not forced to concede that we may pragmatically justify individual beliefs. But the critic thinks that pragmatic justification of entire practices is equally problematic. It seems as though I must admit that we can use pragmatic considerations to justify entire practices, even though we know that these practices are theoretically unjustified. For example, someone might be a doctrinal eliminativist about religious discourse: they might reject the Christian doctrines of sin, salvation, heaven and hell, and so forth.<sup>26</sup> Yet the same person might think that we ought to continue promoting religious belief, and encouraging people to believe the doctrines of the Christian religion. For instance, such belief might (so goes the reasoning) encourage good behavior among people, and provide them with an additional motivation to behave well. Thus, we have an example of a practice about which (a) one is a doctrinal eliminativist, but (b) she thinks we are pragmatically justified in continuing to participate in this practice.

This objection (that pragmatic reasons cannot really justify practices) is right about something: pragmatic reasons cannot license us in believing false things about the world. In other words, pragmatic reasons cannot commit us to objects or causal powers that do not exist. The pragmatic

justification of religious discourse fails *precisely because* it requires that we come to believe something false about the world. That is, if religion is to serve the function of keeping people in line, these people cannot be eliminativists – doctrinal, ontological, or causal – about religious discourse: they must really believe in heaven, hell, God, miraculous powers, and the like if religious discourse is to serve the role envisioned above. But moral discourse doesn't commit us to false beliefs about the world, *because moral discourse need not be a descriptive discourse in the first place*. Moral discourse need not attempt to describe the world, or the moral facts therein. Instead of describing, moral utterances might *legislate* or *endorse*.<sup>27</sup> To issue a normative (e.g., moral) utterance is to endorse our current practice, or to call for its revision. In other words, normative utterances legislate, they need not describe. Thus, if I say "Setting cats on fire is cruel," I am endorsing our practice of condemning hoodlums who set cats on fire. Or, if our practice doesn't currently condemn the igniting of cats, I might say "Setting cats on fire is cruel," thereby calling for the revision of our practice; I am legislating. Coupled with a deflationary account of truth, moral discourse can then be truth-appt without being descriptive.

Pragmatic reasons therefore justify our legislating morally (i.e., making moral claims), and they justify our endorsing and condemning certain actions or practices. But they don't commit us to the existence of any mythical causes or objects – I've argued for this much when I argued for causal explanatory eliminativism and ontological eliminativism about moral discourse. Indeed, moral discourse need not even commit us to the existence of properties such as *cruelty* or *kindness* which inhere in certain actions or practices – such properties might only be the reflection of our moral practice. The practice is prior to these properties, and ascribing these properties to such actions is merely an elliptical way of legislating or endorsing moral rules.<sup>28</sup>

I should express some reservation about this last point. It is not at all clear to me that the pragmatic justification of a practice must commit us to thinking the practice is thoroughly non-descriptive; such a practice might legitimately ascribe *properties* to things. We need only be careful that in using the vocabulary of the practice, we are not thereby committing ourselves to any non-existent *objects* or *causes*. For example, suppose we introduce the term 'plonk' to mean 'gray and metallic.' Thus, my car really is plonk. The predicate 'is plonk' might even be causally explanatory – it might explain why the light glints off of my car in a particular way. As long as use of the predicate "plonk" doesn't commit us to the existence of anything objectionable, there is no reason why we may not pragmatically introduce this new predicate.

Moral practice may well function in a similar way; an action might actually be cruel. If an action causes needless suffering, then it really is

cruel, just as objects that are gray and metallic really are plonk. But unlike religious practice, which cannot be pragmatically licensed because the practice commits us to the existence of God, angels, and mysterious causal powers, the moral practice commits us to no mythical entities or causes. It is merely a pragmatically justified way of dividing up the world; it is a way of dividing actions into those that are to be done and those that are not to be done. As long as we control the consequences of application<sup>29</sup> of any terms that we introduce, a practice may be both descriptive and pragmatically justified. So as long as “P is cruel” doesn’t have as a consequence of application anything like “God exists,” or “Miracles can happen,” then the predicate “is cruel” does not seem problematic. The problem with our pragmatically justified religious practice was precisely that it did commit us to such beliefs. But the properties of our moral practice need be no more problematic than our hypothetical predicate “is plonk”: it is descriptive, it is pragmatically justified, but it commits us to nothing that offends.

At the end of the day, it seems, then, as though moral claims combine both legislative and descriptive elements. To call an action cruel is to characterize it in a certain way, *and* to say that there is at least a *prima facie* reason not to perform this action. Moral discourse, then, is legislative in that its primary purpose is to endorse an action or practice, condemn an action, or call for the revision of a practice. It is also descriptive in that it ascribes properties to actions – it divides actions into those that are to be done, and those that are not to be done. Since this division of our actions into these two categories (and perhaps others – in addition to permissible and forbidden, there are the categories of obligatory and, perhaps, supererogatory) does not commit us to any mythical objects or causal powers, then there is nothing objectionable about this way of dividing up the world.

So pragmatic reasons may not commit us to misdescriptions of the world. They may not be used to justify individual beliefs; nor can they justify us in adopting practices that misdescribe the world. Pragmatic reasons may justify us in pursuing scientific discourse (whose goal is *accurate* description of the world); they may justify us in pursuing moral discourse (whose goal is primarily legislative, and only descriptive in the minimal way outlined); but they cannot justify us in adopting a practice that misdescribes the world, as the above religious discourse would.

Does this mean I have abandoned a commitment to the truth-aptness of moral discourse? By claiming that moral discourse is not primarily descriptive, am I denying that moral claims are candidates for truth and falsity? No; moral claims can be true or false. I merely claim that what it takes for a moral sentence to be true is different from what it takes for a purely descriptive sentence to be true. A descriptive sentence, in order to be true, must be descriptively adequate to the world; it must describe

accurately features of the world. A moral sentence, by contrast, has different truth requirements<sup>30</sup> – say, if it is supported by the best reasons in a practice which is an effective means to advancing certain of our interests – the interests that give rise to our moral practice.<sup>31</sup> So for a purely descriptive (say, scientific sentence) to be true, it must accurately describe the world. If the sentence is, “Electrons have a negative charge,” then for it to be true, the world must contain electrons and these electrons must have a negative charge. On the other hand, for the sentence “Setting cats on fire is cruel” to be true, it need not be the case that it accurately describes anything. Moral sentences must “jump through a different set of hoops” than scientific sentences in order to be counted as true. Thus, if we pragmatically justify our moral practice, we are not committing ourselves to a misdescription of the world, as we would be if we pragmatically justified a set of religious doctrines. Rather, we are committing ourselves to a primarily legislative (and only minimally descriptive) practice.

Is moral truth then a different kind of truth than scientific truth? If the requirements for moral truth are different from the requirements for scientific truth, if moral and scientific truths have to jump through different hoops in order to count as true, then are there multiple kinds of truth? Truth is a large topic, too large for me to do it justice here. For now, I must rest content to show the following: from the fact that morality and science have different “truth requirements,” it does not follow that moral truth and scientific truth are different things. If we have a predicate (say, a truth predicate “T”), the fact that p and q must meet different requirements for T to be predicated of them does not mean that “T” is not univocal. Consider an analogy.<sup>32</sup> Imagine two people, and the requirements they must fulfill to become voters the United States. Margaret was born in the U.S. to American parents; Marguerite, a recent immigrant to the states, was born outside of the U.S. to parents who were not U.S. citizens. Marguerite must jump through different hoops to be allowed to vote; registration will be much harder for her to acquire than it will be for Margaret to acquire. Presumably, Marguerite will first have to become a naturalized citizen, or at least a resident alien (I am unsure of the exact requirements, but this does not affect my argument). They must fulfill different requirements to become voters. But “registered voter” is not equivocal; if both Margaret and Marguerite register to vote, “registered voter” means the same thing in both cases. It is not as though “registered voter” is predicated of Margaret and “registered voter\*” is predicated of Marguerite. *They have achieved the same thing* (registration), even though Marguerite, a foreign citizen, had to fulfill different requirements to achieve it.

So we should not infer from the fact that morality and science have different requirements for counting as true that moral truth and scientific truth are different. They might well be; I do not propose to enter into this

debate. I merely wish to point out that this conclusion does not follow from the mere fact that moral sentences must fulfill one set of requirements to count as true, whereas scientific sentences must fulfill a different set of requirements. Thus, moral sentences can be true; but by committing ourselves to moral practice we are not committing ourselves to a misdescription of the world. This is because morality is primarily a legislative (and only minimally a descriptive) practice; and hence the fact that it is not primarily descriptive is not an objection to it, any more than it is a criticism of a spoon that it does not cut.

### *Moral perception*

We are now in a position to respond to another worry raised by my account. The worry is that if moral facts aren't causally efficacious, then we cannot perceive them, except perhaps with a special sense of moral intuition.<sup>33</sup> Let us briefly discuss why this conclusion is mistaken. Since moral properties do not appear in scientific explanations, it is easy to draw the mistaken inference that they are only perceivable by some special faculty above our five senses. But we need not be driven to this conclusion. Indeed, Harman himself gives us an account of how moral perception might work. We observe actions and institutions, and react to them as we have been trained to do. We have been trained to classify certain actions as cruel, and others as kind. Thus, when we see the hoodlums lighting the cat, we respond as we have been trained to respond: with the belief that the action in question is cruel. Thus, the reaction in question is "due [not] to the actual wrongness of what you see [but is] simply a reflection of your moral 'sense,' a 'sense' that you have acquired perhaps as a result of your upbringing" (Harman 1977). I will concede, then, that moral properties are *non-natural* properties; but this concession amounts only to the claim that moral observations are not the result of causally explanatory moral facts impinging on our perceptual apparatus, but are instead explained by one's moral sensibility. Indeed, although moral properties are non-natural, non-natural properties are properties of natural objects, perception of which is (relatively) unproblematic. We have been taught to classify certain natural items as chairs, others as lions, and others as cruel acts. And while cruelty and warrant may be causally impotent, cruel actions and warranted beliefs are not. Thus, we may causally interact with actions and beliefs. We are trained to group together certain actions as cruel and certain beliefs as justified, but this training applies to all predicates, social and natural. So in the case of both social and natural predicates, a natural object acts on us causally, and we classify it as we have been trained to do so. Hence, there is no more difficulty with moral perception than there is with the perception of

mountains, chairs, stars, or blades of grass. Whether an item is classified as a natural kind or a social kind has no bearing on what faculties are required to perceive it.

### *Morality, epistemology, and reduction*

Harman has another strategy for eliminating moral discourse. Harman perhaps thinks he can pry apart epistemology and morality by arguing that even though both moral and epistemological properties are explanatorily impotent, epistemological properties are reducible to (explanatorily potent) natural properties, while moral properties are not. In this way, Harman hopes to eliminate genuinely normative moral discourse, while retaining normative epistemological discourse.<sup>34</sup>

There is reason to think Harman's strategy will not work. It is important to note that with regard to reduction to natural properties, morality and epistemology stand and fall together. There is no reason in principle why, if epistemological properties are reducible to natural ones, the same cannot be said of moral properties. Naturalism thrives in modern ethical theory much as it does in epistemology. Ethical naturalists abound (see, for example, Brink 1989). A utilitarian, for example, might claim that "morally right" denotes the same property as "maximizes human happiness." Further, the arguments that moral properties cannot be reduced to natural ones will work equally well against any conception of epistemology which conceives of epistemological properties as reducible to natural ones. It is not my purpose to defend ethical naturalism; I merely wish to show that standard objections to naturalistic reduction apply equally to morality and epistemology. Let us very briefly consider a few of these objections.

Consider, for example, an objection discussed by David Brink (1989, pp. 144–70). According to this objection, any attempt to reduce moral properties to natural properties commits the naturalistic fallacy, because moral properties are normative, and natural properties are non-normative. I think this objection is ultimately misguided (and indeed Brink rejects it for the very reason that it conflates identity of concepts with their coextensiveness), but the important thing to note for the purposes of this paper is that this objection could equally be wielded against attempts to reduce epistemological properties (which purport to be normative) to natural (and hence non-normative) properties.

Another version of this objection, also raised by those impressed by the is-ought gap, argues that "the fact . . . that moral claims cannot be deduced from premises that are nonmoral and clearly fact-stating shows that moral claims cannot be fact stating" (Brink 1989, p. 167). But if this is so, then neither can claims about what scientific propositions are justified be

deduced from any set of natural facts, since epistemological claims are normative.

A related objection is G. E. Moore's open question objection (Moore 1903, pp. 6–21). According to this objection, *morally right* cannot be identical to, say, *maximizes utility* because one can say, without logical contradiction, "I know that X maximizes utility, but is it morally right?" It is an open question whether any given action which maximizes utility is also morally right. Again, this objection is misguided (since, as has been pointed out by Kripke (1971) and Putnam (1973), analyticity and necessity are distinct properties); but again, the important thing to note is that this same objection can be wielded against attempts to reduce epistemological properties such as *justified* to natural properties. Consider an example from naturalized epistemology. You can deny that a belief resulting from a certain stimulation of the retina by photons is justified, without thereby logically contradicting yourself. So if the open question objection shows that moral properties are not reducible to natural ones, it shows the same of epistemological properties.

There may be other such objections, but I am confident that the same move I have employed with these three can be employed with the others. There seems no objection to naturalism in morality, which does not also function as an objection to naturalism in epistemology. And in any case, my earlier point about pragmatic reasons stands. If moral (or epistemological) properties are not reducible to natural properties, the pragmatic cost of hewing to the reducibility test of elimination is too high. Why should we regard only scientific practices as worthy of pursuit? Hence, Harman's second strategy for doing away with morality will work no better than his explanatory requirement.

### *Conclusion*

To conclude, then, we started with the claim that morality was eliminable because of its explanatory impotence. I conceded this impotence, but noted that epistemology displayed the same sort of impotence, so epistemological realists and quasi-realists are in no position to demand that moral facts figure in scientific explanations. I then argued that questions of eliminativism cannot be settled *a priori*, but must be settled on a case-by-case basis, and that pragmatic decisions must play a crucial role in the decision. In light of this conclusion, my account of moral truth and the pragmatic justification of moral discourse can be seen as an *a posteriori* justification of moral practice. I conclude, then, that Harman's standard is too rigorous, and ought not be applied in philosophy.

Department of Philosophy  
Auburn University

## NOTES

<sup>1</sup> I am extremely grateful to Mark Lance, Maggie Little, and Henry Richardson, who read numerous earlier drafts of the original paper and provided the crucial feedback that gave the paper its current form. Mark Murphy and Steve Kuhn also read an early draft, and provided excellent constructive criticism. Audiences at the Auburn Philosophical Society and the Alabama Philosophical Society (particularly David Alm, who was commentator at the former event) pressured several weak points, forcing me to further develop and refine my view. An anonymous reviewer for *Pacific Philosophical Quarterly* also provided useful comments.

<sup>2</sup> This line of argument has been advanced in Sayre-McCord (1988). I propose to go into greater detail than Sayre-McCord, and discuss what, if not causally explanatory facts, actually justifies our continued participation in moral and epistemic discourse. The main bulk of the paper is devoted to this latter task: a discussion of the pragmatic reasons that justify both moral and epistemic discourse.

<sup>3</sup> Harman and Ayer are emotivists about moral discourse, a position Harman describes as “moderate nihilism.” Mackie isn’t an eliminativist about morality, but his version of moral discourse is so emasculated that it would be proper to describe it as moderate nihilism, too.

<sup>4</sup> One of the most convincing works demonstrating this point is, of course, Kuhn (1970). Also good are chapter 3 of Longino (1990) and chapter 2 of Churchland (1979).

<sup>5</sup> The terms ‘doctrinal eliminativism,’ ‘explanatory eliminativism,’ and ‘ontological eliminativism’ come from Mark Lance and John O’Leary-Hawthorne (1997).

<sup>6</sup> Some might claim that all purported examples of normative explanation are merely conceptual analysis. If this is true, it does not hurt my position. I am here examining the types of explanations in which moral facts may figure, with an eye to conceding that moral facts do not figure in any interesting way in causal explanations. Since normative explanations are not causal anyhow, a denial that there are genuine normative explanations (as opposed to mere conceptual analysis) is compatible with the concession I am preparing to make to Harman.

<sup>7</sup> Crispin Wright (1992) thinks that an important test for realism is whether a particular discourse enjoys *wide cosmological role* – that is, whether the states of affairs to which the discourse refers may legitimately participate in a wide range of different kinds of explanations. Morality, as we have seen, delivers certain sorts of explanations, but (Wright argues) it does not deliver certain classes of causal explanations. Thus, while moral discourse remains “minimally truth-apt” (i.e., its sentences are candidates for truth and falsity, and have genuine assertoric content), it is not as robustly “real” as, say, scientific discourse, whose referent states of affairs may legitimately participate in a whole range of explanations, from rational to causal. I think Wright is correct; morality doesn’t live up to the standards of science. I only wish to point out that epistemology suffers from the same narrow cosmological role as morality; and I further wish to show why we might want to continue participating in forms of discourse which do not display wide cosmological role.

<sup>8</sup> Even if you deny this – even if you think causal relations are causally efficacious – you do not necessarily succeed in prying apart morality and epistemology. Toward the end of the paper (in the section titled “Morality, Epistemology, and Reduction”) I will argue that if epistemology can be naturalized, there is no in-principle reason why morality cannot be naturalized as well.

<sup>9</sup> I borrow this example from Quinn (1986), who in turn borrowed it from Bernard De Voto.



<sup>10</sup> Some, like Harman (1977, Chapter 2), argue that non-entities such as *average American mothers* are ineliminable because they are reducible to objects which do figure in causal explanations. I will return to this point in the section "Morality and Reduction."

<sup>11</sup> The term 'Average American mother' is not completely analogous to moral terms, because it does not function as a singular term. That is, it can only be substituted into narrowly defined contexts; you cannot, for example, substitute it into the sentence, "Where did you see the X?" However, even if it is not just like a moral term, it does show that ontological and causal explanatory eliminativism do not entail doctrinal eliminativism.

<sup>12</sup> The belief that causal explanatory eliminativism entails doctrinal eliminativism can have particularly nasty effects if one believes that only natural kinds are involved in genuine causal explanations. For if this were the case, then all sentences about artifacts – chairs, shirts, computers, and so forth – would be false. This would call for massive revision of our language, and massive inconvenience – imagine if we couldn't refer to any artifact whatsoever. We will see later on the relevance of such pragmatic reasons in the debate over doctrinal eliminativism.

<sup>13</sup> See chapters 3 and 4 of Kuhn (1970).

<sup>14</sup> Brandom (1994), for example, argues that normative utterances (logical, epistemic, semantic, moral, and so forth) make explicit norms which are already implicit in the practice. Lance and O'Leary-Hawthorne (1997), alternatively, argue that such utterances serve either to endorse the norms implicit in the practice, or to try to bind the members of the practice to a new set of norms.

<sup>15</sup> I am not suggesting that we might be able to infer someone's epistemological commitments entirely from her non-linguistic behavior. I am merely saying that these commitments are carried out in practice, and that this practice is a manifestation (even if a non-interpret-able one) of these commitments. Further, I should add (to avoid rousing the ire of the radical interpreters) that I am not here claiming that we could have epistemic commitments that could not be discovered by a radical interpreter. I am merely denying that non-linguistic behavior alone will always be sufficient evidence for the translator to determine what these commitments are.

<sup>16</sup> As an aside, I should note that pragmatic reasons play an important role *within* science, as well. Every list of criteria for theory choice contains such pragmatic criteria as simplicity and economy.

<sup>17</sup> I owe this example to Mark Lance.

<sup>18</sup> Some have argued the contrary, but it is beyond the scope of this paper to argue that we are better off playing the moral game. Thus, I will have to assume that we are.

<sup>19</sup> I owe this phrase to Duncan Macintosh.

<sup>20</sup> Keep in mind that as a "de-cynicized Hobbesian," I am using 'prudential' in a non-egoistic sense.

<sup>21</sup> Although the case of science suggests an important point: the fact that we pursue a type of discourse for pragmatic reasons does not entail that the discourse must be internally pragmatically structured. After all, the question of whether a particular scientific inference or belief is rational is a question to be answered from within science, according to scientific reasons, and not on pragmatic grounds. I discuss this topic in greater detail in the section "Worries about Pragmatic Justification."

<sup>22</sup> I am grateful to Mark Lance for helpful conversations on this topic. Indeed, the example that follows was suggested by him.

<sup>23</sup> And, of course, instrumentalists will deny that truth is a valid goal in the first place.

<sup>24</sup> This section owes its existence to criticisms raised by David Alm.

<sup>25</sup> From comments by David Alm on an earlier draft of this paper presented at the October 22, 1999, meeting of the Auburn Philosophical Society.

<sup>26</sup> This example was also suggested by David Alm.

<sup>27</sup> Nothing short of a book-length treatment could make good on this claim, so I cannot defend it at length here. Fortunately, someone else has already made progress toward making good on this claim, so I may defer to them. Thus, the following account borrows heavily from Mark Lance and John O'Leary-Hawthorne (1997). I am here offering not an argument for the position advanced by Lance and O'Leary-Hawthorne; I am not really even offering an account of their position. Those who want either should read their book, especially chapter 3 (where the main work of describing the position is done). Further, Lance and O'Leary-Hawthorne do not endorse all of the claims I now go on to make, or all of the uses I make of their work.

<sup>28</sup> I have noted that much of what I am saying here is borrowed from Lance and O'Leary-Hawthorne (1997). However, this claim about properties is a claim which Mark Lance has explicitly repudiated (conversation with Mark Lance). To reiterate what I stated in my previous footnote, it should not be inferred that Lance and O'Leary-Hawthorne would endorse all my uses of their work.

<sup>29</sup> I am here using Dummett's useful term. The *criteria* of application for a sentence are those conditions which license utterance of that sentence. For example, both 'A' and 'B' are criteria of application for the sentence 'A v B.' Once we license utterance of a sentence, we become entitled to that sentence's *consequences* of application. For example, 'A v B' has, among its consequences of application, ' $\sim(\sim A \ \& \ \sim B)$ .'

<sup>30</sup> This talk of truth requirements should not mislead the reader into thinking that I believe we can give robust truth-conditions for moral discourse. It seems likely that sentences of the form "'P' is true if and only if P" tell us about the role of the truth predicate in our language, and not much about what makes 'P' true. My talk of truth requirements in this paragraph and the next is intended only to bring out the fact that very different kinds of discourse can be true; and it doesn't follow from the fact that these discourses are different (say, one is primarily legislative and one is primarily descriptive) that different types of truth are at work in each.

<sup>31</sup> One account of moral truth in terms of excellence of reasons is offered in Wiggins (1991). Social practice accounts of morality must always be on the lookout for relativism. A social practice account of normative discourse that avoids relativism (at least the pernicious forms of relativism) is offered in chapter 3 Lance and O'Leary-Hawthorne (1997). The latter account is unique in that it argues (a) that we cannot give robust truth conditions for moral discourse, and (b) moral discourse is nevertheless truth-apt.

<sup>32</sup> This analogy was suggested in a different context by Roderick Long. You should not infer from its presence that Prof. Long endorses anything I have said in this paper (or elsewhere, for that matter).

<sup>33</sup> This worry is similar to John Mackie's charge that moral properties are epistemically queer (Mackie 1977, pp. 38–42).

<sup>34</sup> I'm not sure that Harman ever suggests this manner of prying apart morality and epistemology, but it is the only view consistent with his corpus. In chapter 1 of (1977), he argues that moral facts do not explain our observations. He then suggests in chapter 2 that a candidate for facthood must either serve such an explanatory role, or be reducible to something which serves such an explanatory role, in order to count as a genuine fact. Harman concludes chapter 2 by remarking that "[s]ince moral facts seem to be neither precisely reducible nor useful even in practice in our explanations of observations, it remains problematic whether we have any reason to suppose that there are any moral facts" (p. 23). As I have noted, epistemological facts display explanatory impotence, as well. So since Harman is a realist about normative epistemological discourse, he must think that the difference between morality and epistemology is that the latter reduces to natural facts,

whereas the former does not. I am ignoring the question of why reducibility to natural facts is relevant to the issue of doctrinal eliminativism. As Quinn (1986) argues, it is not at all obvious how reducibility is supposed to defend against eliminativism. Rather, Harman's caveat here seems a bit *ad hoc*.

## REFERENCES

- Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Churchland, P. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.
- Grover, D. L., J. Camp and N. D. Belnap, Jr. (1975). "A Prosentential Theory of Truth," *Philosophical Studies*, 27.
- Hare, R. M. (1952). *The Language of Morals*. New York: Oxford University Press.
- Harman, G. (1977). *The Nature of Morality*. New York: Oxford University Press. Chapter 1 reprinted as Chapter 6 in G. Sayre-McCord (ed.) (1988) *Essays on Moral Realism*. Ithaca: Cornell University Press, pp. 119–24.
- Kripke, S. (1971). "Identity and Necessity," in S. Schwartz (ed.) (1977) *Naming, Necessity, and Natural Kinds*. Ithaca, New York: Cornell University Press.
- Kuhn, T. S. (1959). "The Essential Tension," in C. W. Taylor (ed.), *The Third (1959) University of Utah Research Conference on the Identification of Scientific Talent*. Salt Lake City: University of Utah Press, pp. 162–74. Reprinted in Kuhn (1977), pp. 225–39.
- (1970). *The Structure of Scientific Revolutions*, 2nd edition. Chicago: University of Chicago Press.
- Lance, M. N. and J. O'Leary-Hawthorne (1997). *The Grammar of Meaning: An Exploration of the Normative Character of Semantic Discourse*. New York: Cambridge University Press.
- Longino, H. E. (1990). *Science as Social Knowledge*. Princeton, NJ: Princeton University Press.
- Putnam, H. (1973). "Meaning and Reference," in S. Schwartz (ed.) (1977) *Naming, Necessity, and Natural Kinds*. Ithaca, New York: Cornell University Press.
- Quinn, W. S. (1986). "Truth and Explanation in Ethics", *Ethics*, 96:3, pp. 524–44.
- Rawls, J. (1955). "Two Concepts of a Rule," *The Philosophical Review*, 64:1, pp. 3–32.
- Sayre-McCord, G. (1988). "Moral Theory and Explanatory Impotence," in G. Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca, NY: Cornell University Press.
- Sellars, W. (1953). "A Semantical Solution of the Mind-Body Problem," *Methodos*, 5, pp. 45–82. Reprinted in J. Sicha (ed.) (1980) *Pure Pragmatics and Possible Worlds: The Early Essays of Wilfrid Sellars*. Atascadero, CA: Ridgeview Publishing Company, pp. 219–56.
- Wiggins, D. (1991). "Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs," *Proceedings of the Aristotelian Society* 91, pp. 61–85.
- Wright, C. (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.