

Moral Value Property Projectivism

Ralph W. Clark

Abstract

I describe and then raise objections to what may be referred to as the Hume/Blackburn account of a projectivist view of the nature of moral values. One of the objections is that a category mistake is present in the claim that attitudes or feelings can be projected onto the world as a factor in the constitution of value properties. In order to avoid this and other objections, “moral value property projectivism” involves the projection of properties rather than attitudes or feelings. These properties are located in individual states of human consciousness. I focus on two such properties, which I designate goodness and badness, and which I argue are most immediately projected in such fashion as to give rise to general value properties. The constitution of moral value properties requires additional steps. The view I defend is a version of naturalism and moral realism.

1. Introduction

What are moral values? Answers can be placed in two broad categories that pertain to what may be called the “location” at a fundamental level for moral values. As regards the first location, in seeking an account of moral values we look primarily to the objects (actions, persons, intentions, outcomes, etc.) about which moral judgments are made. As regards the second location, we look primarily to the subjects (human beings, either individually or collectively) who make moral judgements, or in some other way respond evaluatively to objects. The first location is problematic because the properties of objects – at least such properties as science investigates, namely natural properties – seem bereft of some essential elements, such as the action-guiding character of moral values. Yet, if objects are said – à la Moore – to possess “nonnatural” properties linked to these essential elements, but which lie outside the purview of science, then these alleged properties raise questions that most philosophers would agree have not been satisfactorily answered. Because nonnatural properties by definition elude scientific investigation but are “out there in the world,” yet are causally inert, how can we gain knowledge of them?

The second location is problematic as regards each of the two general types of answers that belong to it – cognitivist and noncognitivist. For cognitivist answers, the most widely defended claim is that moral values are best understood in terms of morally relevant facts about people’s attitudes or desires. The main problem here is that describing attitudes or desires that people happen to have – saying, for example, that a certain group of people are in possession of attitudes opposed to a particular sort of inequality – does not seem to capture fully the action-guiding character of moral value assertions since it does not explain why members of this group, or anyone else, ought to have the attitudes in question.

For noncognitivist answers, the basic claim is that discourse regarding moral values is best understood as expressing in a nondescriptive fashion morally relevant attitudes, desires, or commitments. One of the main problems here is that moral discourse seems to be truth-apt, while the nondescriptive expression of attitudes or desires is not truth-apt, or at least is not straightforwardly or robustly truth-apt.

Noncognitivist positions are standardly described as “expressivist,” but this label can be misleading since some of the advocates of expressivist views deny that their views are noncognitivist. An example is Simon Blackburn, who argues that morally relevant attitudes or desires are “projected” onto that to which moral discourse is addressed. To complicate the picture even further, it is possible – indeed desirable, as I argue in this paper – to defend a projectivist position that is cognitivist but does not have the “expressivist basis” that Blackburn gives to his version of projectivism.

The major answers and types of answers to which I have just referred, the major objections to these answers, and many of the responses to the objections are well-known in the literature. I spell out some of these answers, objections, and responses in what follows, in the course of defending what I believe to be a new version of the second broad type of answer (where we look primarily to human subjects in giving an account of moral values). The position that I defend is a version of moral realism (and thus also cognitivism), understood to be the view that judgments about moral values are, or are the contents of, assertions that are true or false – or at least are not automatically excluded from being true or false. Blackburn describes his position as being an example of quasi-realism. I will refer in what follows to some of the bases for debate as to exactly which senses of “true/false” are available, respectively, to the advocates of moral quasi-realism and the advocates of moral realism.

I begin by discussing the views of Blackburn, who is the best known current advocate of projectivism. I then defend a new version of projectivism that differs in some important respects from the position of Blackburn.

1. Blackburn's Projectivism

The core of Blackburn's position derives from a view made famous by Hume. Consider the following passage from Blackburn (1981, 470):

... in addition to judging the states of affairs the world contains, we may react to them. We form habits; we become committed to patterns of inference; we become affected, and form desires, attitudes, and sentiments. Such a reaction is 'spread on' the world, as Hume puts it in the *Treatise*, by talking and thinking as though the world contains states of affairs answering to such reactions.

Blackburn expresses equal approval for a passage from Hume's *Enquiry* where Hume famously describes the precursor to the projectivist position of Blackburn and others as involving "... gilding or staining all natural objects with the colours, borrowed from internal sentiment ...". As these quotations make clear, the core view of Hume and Blackburn is a variation on the "second location" view that I described above, according to which, in seeking a location for moral values, we look primarily to the subjects who make moral judgments or in some other way respond evaluatively to objects. Blackburn's version of projectivism is also a version of quasi-realism, as I have already mentioned, the intent of which is to show that projectivism, which for Blackburn begins life as a version of expressivism and thus noncognitivism, can be transformed into – or, perhaps better, can be seen as — a view that is not exactly noncognitivist. Below, I say more about Blackburn's quasi-realism. Before doing that, I will begin to present my case against Blackburn's position by arguing that its Humean foundation appears to be incoherent.

The category mistake objection to Hume/Blackburn. To assert that an object – which could be a state of affairs – has moral value is (or certainly seems to be) to ascribe a property to the object, or in other words to make a claim about *the way the object is*. By contrast, to have an attitude toward, a desire for, or a sentiment about an object is *to stand in a relation to the object*. Thus, to speak of "spreading" a sentiment on the world, or "gilding or staining" an object with a sentiment, or an aspect of a sentiment, is,

obviously, to speak metaphorically, but not in a way that advances the cause of philosophical clarification or argumentation. To express an attitude, desire, or sentiment is not to make any claim about the way an object is, or to make any claim at all for that matter, but rather to use language nondescriptively to convey information about the aforementioned relationship – to induce the listener to come to understand that, for example, the speaker favors the object.

The use of metaphor to which I refer in the previous paragraph has all of the advantages of theft over honest labor since it describes one thing (an attitude, desire, or sentiment—none of which is a property) in terms of another thing (namely something that is a property, or a quality; to become gilded or stained is to be changed qualitatively) in a situation where the heart of the relevant metaethical debate concerns the very question of whether or not moral discourse concerns the possession of a moral property by, for example, a state of affairs that is said to be just. If we say that a moral property is present, then we have a solid foundation for defending moral realism; if we deny that a moral property is present but say instead that there is present only an attitude, desire, or sentiment on the part of those who respond to the state of affairs in question, then we no longer have a solid foundation for defending moral realism. What the Humean projectivist metaphor, or metaphors, do is meld together the opposing sides in the debate. This would seem to be an illicit use of metaphor since, if the metaphor is allowed to stand as conveying philosophical content, then we will have gained a way to speak of moral properties from an expressivist perspective without having given a philosophical account of how it is that expressivism allows for “moral property talk.” On its face, certainly, expressivism does not allow for moral property talk.

But perhaps there is a mechanism by which an attitude, desire, or sentiment can be transformed into a property. After all, not all projectivist positions are illegitimate. Lockean secondary quality projectivism, to which the moral projectivism of Hume/Blackburn is sometimes compared, does not suffer from the problem to which I just referred. For the Lockean view, it is said that phenomenal greenness, as an example, is projected onto the leaves of a tree at which I am looking; phenomenal greenness is projected out from my consciousness. While phenomenal greenness cannot exactly be spoken of as being a property which belongs to the conscious state that is present when I look at the leaves of the tree – to describe a conscious state as being green is itself a category mistake – nevertheless, phenomenal greenness is experience-

ed by me as being a “way that something is, or appears,” which is to say that phenomenal greenness is either a property or at least is property-like. Therefore, in order to say that phenomenal greenness is projected onto the leaves, we do not need to employ the objectionable metaphorical leap that I described above.

Someone may argue that a color property such as greenness is not the best type of example to use in drawing a comparison between the projectivism of Hume/Blackburn and Lockean secondary quality projectivism because such a comparison does not involve attitudes or desires. A better example is the property of being delicious. The Lockean basis for saying that an item of food is delicious has got to be an attitude or desire, or some other type of evaluative response. What else could it be? But then it must be the case that the Lockean projectivist view (which for Locke applies to all five of the senses, including taste), or a view that closely resembles the Lockean view, *is* inclusive enough for moral property projectivism to fall within its scope. Lockean secondary quality projectivism, we have noted, does not suffer from the problem that I have been discussing, a problem that I have claimed does apply to the Hume/Blackburn version of projectivism.

My response is that, while the Lockean projectivist view does apply to all five of the senses, including taste, the examples pertaining to taste that Locke himself gives, such as sweetness, are not, strictly speaking, evaluative but instead are qualitative in a fashion that at least resembles the qualitative character of greenness. Just as phenomenal greenness is experienced as being a “way that something is, or appears,” so also phenomenal sweetness is experienced as being a “way that something is, or appears.” Phenomenal sweetness is either a property or is property-like. Thus, in saying that phenomenal sweetness is projected onto an item of food, we do not need to employ the objectionable metaphorical leap discussed above. By contrast, the property of being delicious is evaluative; there does not appear to be a qualitative phenomenal component of it that can be directly compared to phenomenal sweetness. Instead, what we have is a positive attitude on the part of someone who likes the food in question. An appeal to Lockean projectivism does not help in explaining how this attitude, which is neither a property nor property-like, is transformed via projection into a property.

Below, I defend a version of projectivism for moral properties that *is* close enough to Lockean secondary quality projectivism that, like the Lockean view, it does not need to employ the sort of objectionable metaphorical leap that I have been discussing.

In the meantime, I should mention that one way to reply to what I am calling the “category mistake objection” to the projectivism of Hume and Blackburn is to deny that the outcome of this projectivist picture is the claim that that onto which an attitude, say, is projected thereby acquires a property, or more specifically a moral property. However, if this reply is the direction that one goes in, then the resemblance between the moral projectivism of Hume and Blackburn, on the one hand, and the secondary quality projectivism of Locke, on the other hand, becomes even more tenuous than in the above discussion. Why, then, even refer to the former as projectivism? More to the point, regardless of what we may decide to call the view of Hume and Blackburn, for Blackburn there is a reason why he needs to preserve a link to the Lockean position, namely that Blackburn wants to claim that moral assertions are truth-apt. How can an assertion such as “Act A is moral” be truth-apt if “moral” as it occurs in this assertion does not name a property? If we back away entirely from the claim that “moral” names a property, then it would seem that we have retreated to the emotivist type of expressivism of Ayer and Stevenson, which is a position that many current philosophers view as discredited, and which, in any event, is a position that Blackburn himself wishes to distinguish from his own.

The normative arbitrariness objection to Hume/Blackburn. The root error that this objection capitalizes on can be traced back at least to Hobbes’ claim that the definition of “good” is “object of desire” – which places the cart before the horse. Why would anyone ever desire something unless that thing was already good, or judged (or perceived) to be good, before the act of desiring came on the scene? Similarly, why would anyone have a positive attitude toward, or positive sentiment for, an object unless that object was already good, or judged (or perceived) to be good, before the attitude or sentiment came on the scene?¹ Indeed, it is difficult to understand how a desire (or positive attitude or sentiment) would even be present in a situation where no object was present that had, or was seen as having, a value property. After all, everyone agrees that desires motivate. If I desire X, then

¹ A sophisticated version of a Hobbes-like position is found in Gibbard (1994): “To be good is to be desirable, and a thing is desirable if desiring it is warranted.... Non-cognitivists [such as Gibbard himself]... try to explain away the appearance of non-natural properties.... The expressivist’s strategy is to [say] what state of mind is expressed by ascriptions of warrant.... This state of mind ... consists in *accepting norms* that say to do something.” We can ask the same question of Gibbard that we ask of Hobbes: Why accept a norm that tells us to do something unless that something is good (or leads to what is good)?

I am motivated to gain or keep X unless a stronger, conflicting motivation on my part is present. How can I be motivated to gain or keep that which is, or is perceived to be, value neutral? Perhaps I can be thus “motivated” if “motivation” is understood to be wholly dissociated from deliberation or choice, and instead is no more than observed “movement toward” an object.² At best, under such circumstances, an act of desiring that underlies motivation is arbitrary. But surely, not all desires are arbitrary. At worst, acts of desiring that are claimed to be wholly arbitrary are not psychologically possible; they are discriminatory responses that occur in situations where there is no recognized basis for the discrimination.

Analogously, if the moral goodness of a state of affairs is said to be derived from – rather than being a basis for – the act, or acts, of desiring that state of affairs on the part of some person or group of persons, then the ascription of such moral goodness is arbitrary since the same person or persons could just as well have desired something quite different without there having been any change in normatively relevant factors pertaining to the state of affairs in question. For example, consider the moral value that resides in a person’s having a morally praiseworthy character. We want to be able to say that someone has a morally praiseworthy character because – at least in part – that person consistently desires what is morally good or right, while someone else lacks such a morally praiseworthy character because, at least in part, this other person does not consistently desire what is morally good or right. In other words, we presuppose in these situations that the possession of the property of moral goodness by that which is desired precedes the desire. This does not seem to be possible for the projectivism of Hume and Blackburn.

The truth-aptness objection to Hume/Blackburn. This objection to Blackburn’s projectivism is from the perspective of ordinary ways of thinking and speaking. I have already mentioned that Blackburn claims that moral assertions are truth-apt, while also acknowledging that his projectivist account of moral assertions is a version of expressivism. Expressivism, as noted above, is normally classified as falling within the purview of

² This does seem to be Hobbes’ view: “There be in animals, two sorts of *motions* peculiar to them: one called *vital* ... such as are the course of the *blood*, the *pulse* ... the other is *animal motion* ... as to *go*, to *speak*.... These small beginnings of motion, within the body of man, before they appear in walking, speaking, striking, and other visible actions, are commonly called *endeavor* ... This endeavor, when it is toward something that causes it, is called *appetite* or *desire*.... (from *Leviathan*, excerpted in Johnson and Reath (2007), 137–38)

noncognitivism, according to which discourse regarding moral values is to be understood as expressing in a nondescriptive fashion morally relevant attitudes, desires, or commitments – thus entailing that moral utterances are not truth-apt.

The objection that I am addressing here is an objection to which Blackburn himself has responded at length. In a review article, commenting on his own position, Blackburn (2006, 154) says the following:

Quasi-realism was explained as trying to earn, on an expressivist basis, the features that tempt people to realism. In other words, it suggests that the realistic surface of the [moral] discourse does not have to be jettisoned. It can be explained and defended even by expressivists. Perhaps surprisingly, thoughts about fallibility, objectivity, independence, knowledge, and rationality, as well as truth and falsity themselves, would be available even to people thinking of themselves as anti-realists.

How is this to be accomplished? Michael Ridge (2006, 647) succinctly states Blackburn's answer:

Blackburn's story ... is to give a deflationist account of truth and truth-aptness, according to which (*very* roughly) there is no robust property of truth and there is no real difference between saying 'p' and saying 'p is true.' On this way of thinking about truth, the fact that ethical sentences are truth-apt may well not reveal anything deep about the states of mind those sentences are conventionally used to express....

The main problem with this deflationary truth strategy as employed by Blackburn is that it begs the question. For mainstream advocates of the deflationary theory – that is, philosophers working mainly within the philosophy of language and metaphysics – the claim that there is no real difference between asserting “p” and asserting “p is true” is meant to apply to those situations, and only to those situations, where we already have good reason to suppose that asserting “p is true” is appropriate, as when we say, to use a textbook example, that “‘Snow is white’ is true” asserts no more than does “Snow is white.” However, if the assertions on which we are focusing consist of the expression of attitudes, then we would not seem to have good reason to suppose that these assertions are such that assigning truth to them is

appropriate. Of course, people all the time *do* suppose that moral assertions are truth-apt; it is precisely this supposition from ordinary discourse that Blackburn wants to accommodate via his projectivist, quasi-realist account of moral discourse. But Blackburn cannot point to this supposition from ordinary discourse in order to avoid the charge of question-begging without having already shown that ordinary discourse, in addition to its presupposition that moral assertions are truth-apt, also presupposes that expressivism is correct, rather than presupposing that moral realism is correct. No worthwhile philosophical purpose is served by substituting one apparent instance of question-begging for another. Likewise, if we are willing to stipulate that the processes involved in moral discourse projectivism whereby – as Hume describes them, and Blackburn concurs – an attitude is “spread on” an object do in fact yield a moral property as belonging to the object, or at least something that is like a moral property, then again we would have good reason to suppose that asserting “p is true” in a moral context is appropriate – and we could then move on to the task of determining how a deflationary account of truth fits into this picture. Above, I argue that we should not accept the outcome of the processes involved in moral discourse projectivism that I have just described (processes said to yield moral properties as belonging to objects). Therefore, if I am on the right track here, there does not seem to be any way for Blackburn to show that his deflationist truth strategy applied to moral assertions is not question-begging.

2. A New Version of Projectivism; Some Preliminaries

I am giving the name “value property projectivism” to the basic view that I wish to defend in this paper. Unlike other versions of projectivism in metaethics, value property projectivism, as its name suggests, involves the projection of properties rather than attitudes, desires, or sentiments. Value property projectivism thus avoids the problems that attend upon the claim that attitudes, desires, or sentiments can be projected.

Necessarily, since projectivism of any type involves processes that must begin with human subjects insofar as they have experience of objects, if value properties are to be projected, they must primarily be located in, or primarily belong to, human subjects, and more specifically must primarily be located in individual states of consciousness – just as phenomenal greenness is located in individual states of consciousness. There are, I believe, two such value properties that are located in, or belong to, individual states of consciousness,

namely goodness and badness. When an individual state of consciousness possesses the property goodness, then that state of consciousness is experienced to *be good* by the person whose state of consciousness it is. When an individual state of consciousness possesses the property badness, then that state of consciousness is experienced to *be bad* by the person in question. Such instances of goodness and badness do not just by themselves have implications for moral values, the nature of which is more complex. I argue below that moral values (and moral disvalues) are “built up from” such instances of goodness and badness in conjunction with other factors; in themselves, such instances of goodness and badness are nonmoral value properties.

As an illustration, consider the state of consciousness that belongs to someone – let us call him Jim – during the time when this person is enjoying lunch. Jim’s state of consciousness possesses the nonmoral property of being good but Jim is not likely to be aware of this fact, at least not as something on which he is focusing. While Jim’s state of consciousness is at least partially accessible to him via introspection, ordinarily someone who is enjoying lunch would not engage in such introspection but instead would focus on his food or on some other activity that the person happened to be engaging in simultaneously with eating lunch. If asked about the quality of his food or how his lunch was going, Jim would, we are supposing, say something to the effect that the food was good, but would almost certainly not comment on any of the properties of his state of consciousness.

Now, something can be good for many different reasons. I want to focus on Jim’s reasons for saying, if asked, that his lunchtime situation is good *here and now*, as opposed to its being good in relation to future benefits from, say, eating properly or spending money wisely. I submit that the fact that Jim’s state of consciousness possesses the property of goodness underlies anything that he might say to the effect that his lunch is good here and now. As should be apparent, in claiming this I do not wish to comment on the meaning of “good” as used by Jim when he says that the food is good. Instead, I wish to say something about what, most fundamentally, *is good* in Jim’s lunchtime situation as I am describing it.³ Most fundamentally, it is Jim’s state of consciousness that possesses the property of nonmoral goodness. I wish to say further that this property is a natural property but is not – at least at the present time – analyzable. It is a natural property because (1) while at

³ The point that I am making here, and related points elsewhere, are intended to be in accord with what is described as “new wave moral semantics” by Horgan and Timmons (1992).

present it is discoverable only via experience of one's own inner states, where its standing as unanalyzable is not challenged, given an appropriate theory of mind there is no reason in principle why it could not be analyzed in terms of brain states or brain functions; (2) its existence can be invoked as part of a causal explanation of, for example, the fact that Jim desires his lunch.

Question: Should we not say that Jim's state of consciousness while eating his lunch is *good to Jim* in place of saying that his state of consciousness is good *simpliciter*? After all, no one but Jim could ever directly experience his state of consciousness as being good, or directly experience it in any other way for that matter (assuming as correct the widely held view that one's own states of consciousness are private in a relevant sense). What this shows, however, is that the distinction between *good to Jim* and *good simpliciter* does not arise regarding value properties ascribed to Jim's state of consciousness. A somewhat different but related problem does arise regarding the transition from the projection of nonmoral values to the projection of moral values: Are moral values relative to persons or groups in the sense of applying only *to* those persons or groups? I discuss this issue in a later section.

The position that I am defending here can fruitfully be compared to the position defended by Derek Parfit in "Normativity" and elsewhere even though Parfit defends a version of nonnaturalism and locates normative properties differently than I do. What Parfit's position and the position that I am defending have in common is the claim that normative *properties*, rather than attitudes or desires, reside at a fundamental level of analysis for value claims, including moral claims.⁴ I locate the normative properties in states of consciousness, while Parfit locates them in objects and states of affairs that lie outside of states of consciousness. According to Parfit (2006, 331), "Normative concepts form a fundamental category." I agree. Parfit rejects the claim, as I do, that normative concepts must be restricted to the category

⁴ The position that I am defending can also fruitfully be compared to the sort of objectivist position that is the focal point for the attack famously leveled by Mackie (1977) against what Mackie characterizes as the standard way that moral properties are treated in ordinary discourse: This is the view that moral motivation springs directly and completely from an awareness of moral properties. I am concerned with value properties (and derivatively with moral properties) - which, I wish to say, are such that motivation springs directly from them. For Mackie, such motivating properties, if there were any, would reside in external objects, and for that reason would be queer, he says. I am arguing that such motivating properties belong fundamentally to states of consciousness, and for that reason Mackie's charge of queerness does not apply to them.

of that which is moral; practical decisions, he says, need not involve moral thinking. Parfit says (332): “If we believe in irreducibly normative truths, we are what Korsgaard calls *dogmatic rationalists*.” I accept this way of characterizing both Parfit’s position and my own position. One of the consequences of accepting a “dogmatic rationalist” position in value theory is, as Parfit notes, that there is not a lot that is positive (as opposed to negative – criticizing other views) that one can say in defense of one’s position. What an advocate does, essentially, is to bring to the reader’s or listener’s attention the sorts of instances of normative properties that the advocate believes exist. This is what Parfit does, and this is what I am attempting to do.

Question: Why should we say that the goodness of Jim’s lunchtime situation here and now is located fundamentally in Jim’s state of consciousness?

Consider that if Jim were unconscious he would not be enjoying lunch even if he were somehow – perhaps in his (dreamless) sleep – going through the motions of eating. Consider next that some people seem to lose the capacity to enjoy things while retaining the capacity to go on doing those things. As far as Jim’s lunchtime situation is concerned, this would mean that everything could remain the same as regards what Jim is doing and what, if anything, is being done to him – the same act of eating, the same food, the same atmosphere in the restaurant, the same degree of Jim’s being nourished, etc. The only difference would be a lack of enjoyment on Jim’s part, which accordingly would seem to be a conscious phenomenon. The essential variable, therefore, would seem to be Jim’s state of consciousness. More specifically, we might ask: Is Jim *responding consciously* to his lunchtime situation in such fashion that his experience is good? If the answer is *yes*, then it would seem that we should predicate goodness of Jim’s state of consciousness. Once we have done this, we can predicate goodness in a derivative fashion of other elements in the situation, such as the food Jim is eating or the conversation he is having with a companion. Below, I say more about such derivative applications of the predicate “goodness.”

First, I will address the following question: In order to be good, must Jim’s state of consciousness be pleasurable? Consider that it makes sense to say that Jim’s state of consciousness is good in situations where references to pleasure are not necessary and perhaps are flat-out inappropriate, as for example a quiet evening at home for Jim where his state of mind is primarily relief that a stressful day is coming to an end. Hedonists such as Epicurus

sometimes appeal to the idea of modest or tranquil pleasure that may seem indistinguishable from someone's simply not experiencing pain, in order to include cases where, if ordinary language is our guide, we would probably not want to say that pleasure is present.

In addition to the ordinary language-based objection to the hedonist's attempt to include all instances of intrinsic good within the scope of the term "pleasure," hedonism is vulnerable on the basis of notorious difficulties that are present when any attempt is made via introspection to locate an element that can be characterized as pleasure on any and all occasions when someone can correctly be characterized as having a good, or positive, experience. These introspective difficulties for hedonism, to which I cannot do justice here, have been extensively discussed elsewhere.⁵ Because value property projectivism, as I am defending it, does not appeal to the concept of pleasure, it is in no way compromised by any of these difficulties.

As regards Jim's lunchtime situation, the most important thing to be said in the context of the present discussion is that something about the situation is good. Most fundamentally, I submit, it is Jim's state of consciousness that is good. Generally speaking, nothing besides a person's state of consciousness is universally and noncontroversially present on any and all occasions when we want to say that the person's experience here and now is positive, or good. Moreover, this picture of the situation nicely fits with the link between goodness and motivation. Most straightforwardly, desires motivate because they are oriented toward goodness, as I observed earlier; we might just as well say that goodness motivates. Why, then, need goodness be said to motivate only via the intermediary of pleasure?

The view that I am defending here is consistent with a natural way of describing the phenomenology of motivation by values, moral or otherwise. We do not normally say that we are motivated by our desires (some of which, after all, we do not wish to allow to possess motivating power), but rather by our reasons, which are oriented to the properties that we apprehend desired

⁵ A good treatment is Katz, Leonard. "Pleasure" *The Stanford Encyclopedia of Philosophy* (Fall 2007 Edition), Edward N. Zalta (ed.), <http://plato.stanford.edu/archives/fall2007/entries/pleasure/> See especially 8–13. A good statement of one aspect of the problem is Feldman (2004, 79): "Consider the warm, dry, slightly drowsy feeling of pleasure that you get while sunbathing on a quiet beach. By way of contrast, consider the cool, wet, invigorating feeling of pleasure that you get when drinking some cold, refreshing beer on a hot day. Each of these experiences involves a feeling of pleasure - a sensory pleasure, in my terminology - yet they do not feel at all alike. ... I have come to the conclusion that they have just about nothing in common phenomenologically."

things as possessing. On the basis of such reasons, we may yield to a particular desire as regards motivation. On the basis of other reasons, we may dig in our heels and not yield to some other desire.

Question: If what I have been saying about the possession of goodness by states of consciousness is correct, then why is it that our beliefs do not ordinarily reflect the fact that this is so? Instead, we tend to believe that goodness, in many if not most instances, belongs fundamentally to objects, as opposed to states of ourselves. We believe that food is good. Or, if food is judged to be mainly a means to an end, then health is good or life is good. Jokes that are genuinely funny are good. Knowledge and beauty are good.

By way of answering the question posed at the beginning of the previous paragraph, I ask readers to consider that, in terms of the evolutionary development of human beings, there is little if anything to be gained toward a person's survival from the possession of an inclination to focus at all frequently on one's conscious states, while much is to be gained from the possession of an inclination to focus on the details of, say, what one is eating or the people with whom one is dealing. Knowing about the properties that we experience as being possessed by food will help us to obtain food and guide us when we eat it, and knowing about the abilities, dispositions, intentions, etc. of people whom we encounter will help us to benefit from our encounters. We have, therefore, developed a strong inclination to focus on external objects, not states of our consciousness.

At the same time, evolutionary considerations dictate that there be an inner mechanism that rewards or punishes us for, respectively, engaging with things that have survival value and engaging with things that do not have survival value. An *inner* mechanism is required since rewards and punishments that arise from circumstances outside of a person, such as from the existence of parental or social sanctions for harmful behavior, will not be constantly present, while potential benefits and injuries from the objects one encounters are constantly present. And even more important, there does not seem to be any way to conceive how external sanctions – those residing in circumstances outside of a person's consciousness – would actually function as sanctions except via the role of inner sanctions, namely rewards and punishments that are manifested within a person's consciousness. How, for example, would the threat of imprisonment as punishment function as a deterrent if the person being threatened did not perceive the experience of imprisonment, or some consequence of imprisonment upon future experiences, to be bad? The one constant accompaniment for all of our

experiences of objects in the world is our awareness of those objects, so it makes sense from an evolutionary perspective that rewards and punishments would be built into states of consciousness that are part and parcel with our states of awareness – as long as the proviso is attached that these inner rewards and punishments will play a background role in our experiences so as not to distract us from the all-important need to focus our attention on the objects that we encounter. I submit that the most general and basic way to characterize such inner rewards and punishments is to say that the states of consciousness associated with the rewards and punishments possess the properties of goodness and badness. An understanding of such properties, in conjunction with the evolutionary demand that we focus our attention outwardly yields value property projectivism.

A possible objection: Is it not the case that, even if we “defy evolution” and make a point of using our introspective powers to focus extensively on our states of consciousness rather than external objects, we will come up short in our attempts to focus our attention on the alleged properties of goodness and badness belonging to our states of consciousness? Just as hedonism faces introspective difficulties, will not a search for these alleged properties face introspective difficulties? My response is that value property projectivism does not require introspective success in fixing on the properties of goodness and badness as pertaining to our states of consciousness. Although it seems to me that such properties are accessible to introspection, I am willing to acknowledge that perhaps they are not, and if this should prove to be the most defensible view I would reframe my characterization of them. I would not treat these properties as phenomena open to introspection but instead as phenomena to be hypothesized. I would say that value property projectivism rests upon the hypothesis that these properties exist as belonging to states of consciousness.

In terms of the phenomenology of experiencing the evolutionary theory-oriented rewards and punishments to which I have been referring, it makes sense to say that we project out onto the things in the world certain value properties. Because the projection functions so seamlessly, as I am supposing that it does, we are usually not aware that it is taking place, but rather come to believe that it is objects such as items of food that are good or bad. In what follows, I say more about how projection works with moral values, and I compare projection as it pertains to values with projection as it pertains to secondary qualities such as colors. I discuss the link between projected values and motivation.

3. From Nonmoral Values To Moral Values

How do we get from the projection of nonmoral values to the projection of moral values? Let us begin by looking at the least problematic type of property projection, that involving Lockean secondary qualities.

Regarding color ascription, as I have noted, in the case of greenness it makes sense to say that “something in my mind” (phenomenal greenness) is projected by me – as a psychological act on my part – onto the leaves of the tree at which I am looking. What occurs is a private operation within my consciousness that happens to stand in a relation to an external object. Because other people (all who are normal color observers, like me) carry out essentially the same operation, there is an objective basis for saying that the correct answer to the question, what color are the leaves, is that they are green. Contrast this with how people experience the textbook substance pheno-thio-urea, which only a minority of the human population perceive to be bitter. This minority only, we can say, project phenomenal bitterness onto instances of experienced pheno-thio-urea. At the same time, this minority is large enough that the concept of normal observer has no application as regards human sensory experience of pheno-thio-urea, and as a consequence, there is no objective basis for saying that the correct answer to the question, what is the taste of pheno-thio-urea, is either that this substance is bitter or that it is not bitter. Is there, then, an objective basis for saying that pheno-thio-urea is *bitter to* someone who belongs to the relevant minority of perceivers? The term “objective” is broad enough that an affirmative answer to this question is certainly acceptable. After all, it is a fact that a particular person who belongs to the population in question does experience pheno-thio-urea to be bitter. Moreover, if the human population changed in its makeup so that all or most people came to experience pheno-thio-urea to be bitter, then we (the human population as a whole) would have acquired an objective basis for saying that pheno-thio-urea is bitter simpliciter, just as we can say of the leaves mentioned earlier that they are green simpliciter.

Now, let us look again at the example of Jim who finds his lunch to be good here and now. I have argued that he is projecting the property of goodness onto his lunch, but of course we cannot say that Jim’s lunch is good simpliciter since other people – with as much claim as Jim to being normal observers, or evaluators, of lunch-type foods – may eat the same food that Jim has eaten and not like it. Therefore, the situation is similar to that of pheno-thio-urea, which is *bitter to* some people but not to others.

We have, then, the basis for what may look like common ground between projectivism applied to certain Lockean secondary qualities and projectivism applied to values such as *goodness to P* (where “P” stands for some person or group). However, this common ground does not extend to moral values. Consider, first, that for *nonmoral* values, in order to progress from *goodness to P* to *goodness simpliciter*, all we need is a population with the right composition. Consider the obvious value of relaxation (on appropriate occasions). Relaxation on an appropriate occasion is good to Jim, let us say, but it is also good to virtually every other human being; accordingly, it is correct to say that it is good simpliciter – but only as long as we do not mean that it is morally good, or (if there is reason to prefer “moral rightness” to “moral goodness” as a basic value predicate) that it is morally right to pursue relaxation on the occasions in question. Perhaps one’s moral duty will not be fulfilled if one relaxes. At most – in terms of certain ethical viewpoints – relaxation is only *prima facie* morally good or morally right to pursue. We certainly cannot say without qualification that relaxation is morally good or morally right, while we can say without qualification that relaxation is good in a nonmoral sense of being good – and in this respect relaxation is like a great many other nonmoral objects or states of affairs.

Why the difference between the moral and nonmoral cases? To locate the basis for an answer, let us consider colors again. In the case of greenness, if my color perception is out of sync with that of the large majority of observers (I suffer from a type of red-green color blindness), then we might say, not that the leaves are red, but that they are “red to me.” Correspondingly, if I were a masochist it might be said that certain types of pain are “good to me” (I project goodness onto them). However, “moral to me” makes no sense unless we accept a radical version of individual moral relativism that no actual advocates of metaethical projectivism at the present time, or advocates of any other major metaethical theory for that matter, would accept. Thus, the analogy with colors clearly breaks down at this point, and this is instructive because it indicates that common ground between nonmoral and moral values also breaks down at a certain point. One of the lessons here is that we cannot think of moral values as simply being “built up from” nonmoral values via the mechanism of arriving at population-wide convergence as regards what – without that convergence – would simply be instances of *goodness to P*. Something more is needed.

If the situation were otherwise, moral values would be wholly conventional, the consequence of whatever (initially nonmoral) value

convergence a population happened to experience. For example, starting from the observation that neglecting all considerations of justice is good to Jim (on the supposition that Jim is the sort of person for whom this is true; he projects goodness onto all relevant situations where justice is neglected), we would be able to proceed to the conclusion that neglecting all considerations of justice is morally good simpliciter if the requisite population convergence were to come about or to be discovered (which would encompass a situation where most people in the population turned out to be like Jim, not caring at all about justice), which is an unacceptable conclusion. Whatever else we may want to say about the projection of value properties, we want to be able to say that not all projections of value, even where populations are unanimous in embracing them, can count as projections of moral value.

Hume recognized this problem, and attempted to solve it by stipulating that only “disinterested approbation” qualified as the sort of attitude that, when projected, constituted a moral response. If I have more sympathy for my friend than for a stranger, when both are being treated with the same degree of injustice, then I need to correct for my personal bias, according to the view of Hume, in order to exhibit an appropriately disinterested sympathy if my response is to be considered a moral response. The problem with Hume’s position is that the resulting determination that act A (reflecting disinterested approbation) is moral, while act B (reflecting biased approbation) is not moral, becomes a matter of definition, and thus convention: Ordinary usage – as Hume understands it – assigns “moral” in the one case but not the other. Aside from taking note of the attitudes or sentiments that we may happen to find ourselves and our fellows possessing on a given occasion, we have no more reason to act morally than to act nonmorally – which apparently was Hume’s own view, but not something that he worried about because, for one reason, he apparently believed in the existence of a *de facto* uniformity, or near uniformity, of the human population as regards the possession of disinterested approbation in relevant situations.⁶

We know now that the overall human population is much less homogeneous than Hume imagined. Of course, “morally homogeneous societies” do exist and have existed throughout history, especially within primitive cultures. An examination of such societies will allow us to construct a simplified picture of the transition from nonmoral values to moral

⁶ “The minds of all men are similar in their feelings and operations.” (from the *Treatise*, quoted in Johnson and Reath (2007, 175))

values from the perspective of value property projectivism. After doing that, we can move to a discussion of moral values in our contemporary society, which is not morally homogeneous.

A morally homogeneous society is one in which few, if any, individuals raise serious questions about the moral values of the society. Essentially, the concept of moral progress is absent in such societies: There can be progress regarding compliance with moral norms, but not regarding the norms themselves. (Compare: No sense can be attached to the idea of “color perception progress” in situations where “normal color observer” comes to range over a different population than previously.) Thus, from the perspective of the large majority of people in the morally homogeneous society whom we judge to be moral individuals (virtually every society has criminals and misfits, and among these individuals some bona fide amoralists, all of whom we exclude for present considerations), the Humean picture is satisfactory if we interpret it from the perspective of value property projectivism. A moral individual will, first, project nonmoral values, just as Jim does in the lunchtime situation mentioned above, and will act on these values when they do not conflict with moral values. When there are conflicts, a moral individual will – at least usually or largely – be motivated more strongly by moral considerations. What this means, first, is that a moral individual P projects the property of goodness more strongly onto those specific items (actions, intentions, states of affairs, rules – whatever) to which P’s society assigns moral recognition than onto any other items; moral recognition simply reflects unanimity, or near unanimity, as regards the set of items onto which members of P’s society project goodness. (For Hume’s society, if we suppose that Hume described it correctly, the primary such items were actions undertaken in response to disinterested approbation as projected onto the same set of items by most members of the society.) Another way to describe the situation is to say that P projects *overriding goodness* (which is the strongest version of goodness available to P), onto the items in question – at least usually or largely.

The qualification “at least usually or largely” needs to be added because few individuals – even those whom we would describe as being thoroughly moral – can be expected always to embrace doing the exact morally right thing as determined by the values of the society to which these individuals belong. Therefore, we need to add a second condition that P must meet in order to be described as moral, namely that P projects goodness, or overriding goodness, onto a general inclination on P’s part to conform to his

society's moral values. The possession of this general inclination makes up for P's occasional moral lapses as regards specific actions, while allowing P to be a "thoroughly moral" individual. This second condition involves projecting value properties onto the projecting of value properties. Another way to describe this second condition is to say that P possesses a moral self-image which is a more or less constant element in P's psychological makeup, and serves to guide and prod P even when P may not project goodness or overriding goodness onto specific items to which P's society assigns moral recognition. P thinks of himself as being a moral individual overall, and he is correct in thinking this.

4. Defending Moral Value Property Projectivism (First Round)

I wish to defend the following: For morally homogeneous societies, if we accept the perspective of value property projectivism then there will be no problems or complications in saying that moral assertions are truth-apt, that moral realism, naturalism, and cognitivism are correct, that the perspective of the amoralist can be accommodated, and that moral beliefs motivate (except for the case of amoralists, but they do not pose a problem, as I explain below). For purposes of discussion, let us suppose that for a morally homogeneous society S we have a suitable definition for "just act" and that acting justly is the primary moral value in S. Let us suppose also that an individual member M of S asserts the following:

(1) Acting justly is morally right.

Moral properties are natural properties. "Morally right" is a moral property. It *means* that which one is obligated to do. It *is* that property which is projected onto a class of acts that, for a morally homogeneous society, are determined by the moral orientation of that society. Thus, moral members of the society find themselves projecting goodness, or overriding goodness, onto whatever belongs to the class of acts in question. Just as greenness is a natural property, although a projected one, so also is moral rightness.

(1) *is true.* As understood by M and other members of S, (1) is true: Just acts *are* acts which possess the property of goodness, or overriding goodness, and they are experienced as such by most members of S. Most members of S (because they are moral) will simply find themselves projecting goodness, or overriding goodness, onto a certain set of acts, just as most people (because

they have normal color perception) find themselves projecting phenomenal greenness onto a certain set of objects.

Amoralists are accommodated. Amoral members of S are individuals who do not find themselves projecting goodness, or overriding goodness, onto many, if any, of the appropriate items, but they may nevertheless understand enough about how the term “morally right” is used extensionally by members of S to be able to assent with understanding to the claim that (1) is true. In a somewhat similar fashion, some blind or color-blind people can learn to determine for some occasions that it is appropriate to say that a particular object is green.

Moral realism and cognitivism are accommodated. If value property projectivism for morally homogeneous societies makes moral rightness a natural property, and (1) true, then value property projectivism supports moral realism and cognitivism for morally homogeneous societies.

5. Defending Moral Value Property Projectivism (Second Round)

I now turn to the task of defending value property projectivism for societies that are not morally homogeneous. The crucial difference between societies that are, and societies that are not, morally homogeneous is that the latter, but not the former, allow for genuine moral progress wherein not only can there be improvement as to degree of compliance with existing moral values, but there can also be improvement in the set of moral values against which degrees of compliance are measured.

The key for moral realists who are projectivists is to recognize that projecting value properties is a natural process that necessarily is constrained by the nature of that of which the process is a function – namely, human beings, who have a variety of needs, both individual and social. At any given time, there will be in place on the part of various members of society the projection of goodness, or overriding goodness, onto various objects and states of affairs. Some of these projections will be coordinated with other projections, some will not be coordinated with other projections. An example of situations involving the former could be states of promise-keeping, while an example of situations involving the latter could be states of seizing and eating any food a person happens to encounter, regardless of who owns the food. On the whole, coordinated projections will be the basis for successful actions (actions where the projected value property is realized) more often than uncoordinated projections. For obvious reasons, disinterested approbation (to cite Hume’s category) can be expected to have a relatively high

success rate. A high success rate translates into the perception that being good, or being good in an overriding fashion, as projected, really is good, or good in the strongest sense.

The question, are actions of type A *really* moral, can be seen as a question about the likelihood that actions of type A will in the future prove to be such that they can be coordinated with other actions onto which the members at large of a society are projecting goodness, or overriding goodness. Such a likelihood will necessarily be calculated from the perspective of a given time and location where various value projections are already in place – so it is in this sense context-bound, and has a conventional element. In other words, moral progress is always a progression from what happens to exist at a given time; it is “progress from” what exists as opposed to “progress toward” some already-determined “moral ideal.” For the view that I am defending in this paper, no such moral ideal exists – indeed, there is no way even to conceive what it might mean for there to be such an ideal.

For the view that I am defending, we do not *start* by, say, analyzing human nature a la Aristotle, and then *conclude* that, say, the morally right course of action involves promise-keeping (or some other example of a socially cooperative action). Instead, we start with whatever examples of goodness, or overriding goodness, are in fact being projected; then we predict that a higher level of coordination among projected values will be realized if there is a greater degree of promise-keeping; finally, if we observe that the prediction is borne out, we can then conclude that moral progress lies in the direction of having in place for a society the enhanced degree of promise-keeping. Prior to the time when the prediction is actually borne out, we may have good reason to believe that it will be borne out. And indeed, we do have reason to believe that moral progress would occur if there existed a higher degree of promise-keeping than presently exists within virtually any society; we may conclude, if we wish to do so, that we are thereby provided with an insight into human nature.

Likewise, for the view that I am defending we do not start with the “moral intuition” that there would be moral progress if more people acted with the intention to achieve the “greatest good for the greatest number.” Instead, we may predict that there would be moral progress if more people acted with the intention to achieve the “greatest good for the greatest number.” As it happens, this is a prediction for which there does not appear to be much support, and for several reasons that contemporary normative theorists have noted, such as that the utilitarian criterion is too demanding upon individuals

– which translates into saying that nothing close to unanimity is ever likely to occur within any society as regards the projection of goodness, or overriding goodness, onto actions intended to promote the greatest good for the greatest number. For the view that I am defending, the utilitarian criterion does not allow sufficiently for coordination among all the values that the typical individual living in today's societies does now, and will likely in the foreseeable future, be projecting.

Is there a general moral criterion that might fare better than the utilitarian criterion? I suggest that we specify in a somewhat bland fashion a moral criterion that calls for what can be described as a balance of self-interest against the interests of others (the “Welfare-Balancing Criterion,” or WBC); the main goal in appealing to this criterion is to avoid the objection against utilitarianism that it is too demanding upon individuals. I do not have space enough here to expand WBC so as to spell out how WBC might meet other objections to utilitarianism or to other normative views.

Support for WBC would mean that a prediction is borne out to the effect that, the more that humans know about themselves and the world, and the greater the range of their experiences, the more likely it is for most individual human beings to be more strongly motivated by WBC than by other general moral criteria in situations where these other criteria yield different value projections. Keep in mind that for value property projectivism, motivation (all types, including moral motivation) is purely a matter of the presence within the consciousness of the person who is motivated by an instance of goodness, or overriding goodness. As far as motivation is concerned, these nonmoral properties do all the heavy lifting. We describe the resulting motivation as being moral when one or more properties resulting from projection of these nonmoral properties is a moral property, that is, one involving a high level of coordination as regards value projection for members of a society.

6. Conclusion

The conclusion that I wish to reach in this paper is that there exists a projectivist view of moral value properties that is cognitivist but does not have the problematic “expressivist basis” that Blackburn gives to his version of projectivism. In the view defended here, what is projected is a property, not an attitude, desire, or sentiment. Most fundamentally, it is a value property, which may become a moral value property within the right context. The projection of value properties is sufficiently similar to the situation with

Lockean secondary quality projection that major strengths of the latter carry over to the former. The transition from value property projection to moral value property projection can be accomplished satisfactorily, as regards at least the preliminary presentation of this transition that is given in the present paper.

Acknowledgements

An earlier draft of this paper, presented at a conference arranged for members of the West Virginia University Philosophy Department, benefitted greatly from comments from several of my colleagues. Ted Drange provided many helpful comments at a still earlier stage in my thinking about the issues in this paper. An anonymous referee for this journal provided many helpful comments on the penultimate draft of this paper.

References

- Blackburn, Simon. (1981). Reply: Rule-Following and Moral Realism. In Stephen Holtzman and Christopher Leich, (eds.) *Wittgenstein: To Follow a Rule*, pp. 163–187. London: Routledge. Reprinted in Andrew Fisher and Simon Kirchin, (eds.). *Arguing About Metaphysics*, pp. 470–488. London: Routledge.
- _____. (2006). Antirealist Expressivism and Quasi-Realism. In Copp. (2006), pp. 147–62.
- Copp, David, (ed.). (2006). *The Oxford Handbook of Ethical Theory*. Oxford: Oxford University Press.
- Feldman, Fred. (2004). *Pleasure and the Good Life*. Oxford: Oxford University Press.
- Fisher, Andrew and Kirchin, Simon, (eds.). (2006). *Arguing About Metaethics*. London: Routledge.
- Gibbard, Alan. (1994). Meaning and Normativity, *Philosophical Issues*, Vol. 5, pp. 95–115.
- Horgan, Terence and Timmons, Mark. (1992). Troubles for New Wave Moral Semantics: The Open Question Argument Revived. *Philosophical Papers*, 21, pp. 153–75. Reprinted in Fisher and Kirchin. (2006), pp. 179–199.
- Johnson, Oliver and Reath, Andrews. (2007). *Ethics: Selections From Classical and Contemporary Writers*, (tenth edition). Belmont, CA: Thomson.

Moral Value Property Projectivism

- Katz, Leonard. Pleasure. *The Stanford Encyclopedia of Philosophy*, (Fall 2007 Edition. Edward N. Zalta, (ed.), <http://plato.stanford.edu/archives/fall2007/entries/pleasure/>)
- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin.
- Parfit, Derek. (2006). Normativity. In Shaffer-Landau, (ed.). (2006), pp. 325–380.
- Ridge, Michael. (2006). Saving the Ethical Appearances. *Mind*, Vol. 115, pp. 633–650.
- Sayre-McCord, Geoffrey, (ed.). (1988). *Essays in Moral realism*. Ithaca: Cornell University Press.
- Shaffer-Landau, Russ. (2006). *Oxford Studies in Metaethics*, Vol. 1. Oxford: Oxford University Press.
- Swinburne, R. G. (1976). The Objectivity of Morality. *Philosophy* 51, pp. 5–20.

Ralph W. Clark
West Virginia University
256 Stansbury Hall
Ralph.Clark@mail.wvu.edu