MORAL KNOWLEDGE AND MORAL FACTUALITY

RON WILBURN

Abstract

For naturalistic and non-intuitionistic moral realists, moral knowledge is more problematic than ordinary and scientific factual knowledge. For without special faculties of moral discernment, how could we ever detect moral facts and properties? Physical facts and properties may be accessible to perceptual recognition. But how could moral facts and properties ever be similarly accessible? To address this challenge, we need a meta-ethical account that does two things. First, it must explain how the discernment of moral facts and properties ultimately consists only of the detection of appropriate physical items. Second, it must explain why, despite this fact, moral perception seems so very puzzling. In this paper I endeavor to provide such an account. It is largely because of the relational nature of moral properties, and the corresponding externalistically determined normative content of moral property terms, I argue, that our epistemic access to moral knowledge appears mysterious. The metaphysics of moral factuality does a lot to explain the seeming elusiveness of moral knowledge, and in ways that are surprisingly mundane.
Keywords: Moral knowledge, moral factuality.

Introduction

Let's take moral naturalism to be the view that our emerging scientific picture of the world is roughly accurate as a total causal account of nature, and that moral facts, to exist at all, must accommodate this picture. Let's take moral cognitivism to be the view that moral assertions are comprised of truth-values to which we may gain access naturalistically. Putting these two views together, we can speak of naturalistic cognitivism as the position that moral facts exist in, but supervene upon, natural facts, and that our knowledge of moral facts is explicable without reference to intuitionistic epistemic capacities. So understood, naturalistic cognitivism faces at least two formidable challenges, one metaphysical and the other epistemic. Metaphysically, it must explain how normative facts could ever amount to or emerge from more recognizably natural facts. Epistemically, it must explain how such moral facts could ever be effectively knowable. Although this paper is concerned both with these two challenges and with the connections between them, my primary suggestions concern the latter. Metaphysically, I argue that moral properties are radically relational properties. Epistemically, I argue that it is largely because of the said relational nature of moral properties that the route to moral knowledge may seem particularly troublesome. Thus, by proffering the following account, I aim to perform prescriptive metaphysics toward epistemological ends. The metaphysics of moral factuality, I maintain, does a lot to explain the seeming elusiveness of moral knowledge, and in ways that are surprisingly mundane.

Moral Realism: A Background Account

Because our central concern here is to make a limited epistemological point, we need to

introduce a number of assumptions concerning the type of moral metaphysics that we are presupposing as a background. Most of these assumptions are familiar ones, parts of an amalgam of tenets attributable to various Cornell realists. As I have tried to do elsewhere, I do not attempt here to develop and defend this background theory in any detail (e.g., I do not attempt to defend my provisional examples of virtues or primary goods). Even though this is an important project, it lies well outside the scope of this paper. *By offering the following outline, I aim merely to provide us with a reasonably substantive point of departure from which we might derive some sanguine epistemological morals.*

Although Cornell realism is hardly a settled or univocal meta-ethical position, a number of themes recur throughout its exposition by the likes of Boyd, Brink, Railton, Sturgeon and Platts.[1] Its central premise, that mind-independent moral facts obtain, is typically buttressed by many of the guiding assumptions that have revived *non-moral* metaphysical realism in recent decades, e.g., a thoroughgoing naturalism (in which moral facts consist of nothing more than natural facts), a Kripkean causal theory (eschewing reductive analytic connections between moral and natural words and concepts), and an externalistic take on moral motivation especially accommodative to these two preceding views, on which moral facts drive us to action by virtue of something other than intrinsic, in-dwelling motivational powers.

The variant of Cornell realism presupposed here explicitly uses these and other tenets. One of these additional presuppositions is a seminal premise of Davidsonian semantics. Another stresses the centrality of "thick" moral concepts. Most prominently, however, this account employs a schematic theory of the virtues grounded in the notion of a "homeostatic cluster concept," a notion that we may take to neatly supplement common realist notions of moral supervenience. Let's briefly canvass each of these components in sequence before recounting a story about how they might fit together. This leads to a discussion of the relationalist character of moral properties in Section III, which leads, in turn, to a discussion of externalism-related issues in Section IV. In Section V we are finally in a position to discuss the epistemological upshot with which this paper is primarily concerned.

Of the "additional presuppositions" cited above, the first is the most easily described. It is the recommendation of Davidsonian semantics' seminal tenet, that we should take truth, not meaning or reference, to be basic and unanalyzed.[2] Note that this proposal need not conflict with Kripkean semantics' most essential insight, because the former still allows us to assume that causal relations with reality ultimately underlie our talk. The proposal simply takes these causal relations to stand principally between speakers and "states of affairs" (considered as nothing more than the world's being occurrently such and such), rather than between speakers and objects. On this account, to ask for sentential meaning is to ask for no more than appropriate truth conditions, and to ask for term meaning is to ask for no more than appropriate reference or satisfaction conditions.[3] However, to ask for these things is to ask for no more than the relevant T-sentences and reference and satisfaction axioms of a Davidsonian theory of interpretation. The result is what McDowell and Platts label "austere" semantics.[4] T-sentences give us the literal meanings of the sentences they mention. They communicate sentential content, which can then be manifested in various

ways through the performance of various speech acts. Reference and satisfaction axioms then derivatively secure term meaning, but only at a high level of speculative theoreticity.

When we adopt this strategy of truth-conditional analysis and its concomitant encouragement to employ the full resources of our moral discourse, the second of the above-cited "additional presuppositions" becomes most clearly motivated. This, again, is the assumption that "thick" moral terms (e.g., "loyalty," "courage") should be emphasized no less than "thin" moral terms (e.g., "good," "right") in moral analysis.[5] Little is encoded into the unpacked and platitudinous notion of a "good person." However, much is contained in the concept of "personal loyalty." Whereas the former notion is initially too amorphous to be readily applicable, the latter is much more strictly delimited by public standards. This fact is enhanced by the truth-conditional approach, which makes the meaning of "loyalty" no less substantive than the phenomenon of loyalty itself. The full meaning of "loyalty" is likely to outrun the limited and colloquial understanding we have of it at any given time, just as the truth values of claims about loyalty are likely to outrun the verification conditions we have for assessing them. To emphasize a "thick" moral vocabulary is not to eliminate all reference to "the good" and 'the right" from moral theorizing, however. These thinner concepts' place is secure, but only at a high level of abstraction.

This brings us to the third of our "additional presuppositions," i.e., a schematic picture of the virtues that promises to be particularly amenable to Cornell realists' views of moral supervenience. Simply dispensing with all "thin" moral concepts such as "the good" and "the right" would be undesirable, because it would sabotage any hopes we might ever have to describe relations between particular virtues. Rather, "thin" moral properties are better viewed as consequent upon the harmonious interaction of numerous other "thick" moral properties. Think of Aristotelian accounts of flourishing, which give human excellence its most substantive characterization in terms of the harmonious interactive systematization of lower-level, mean-enforcing, desire-regulating character traits. Such accounts construe human excellence only provisionally as an activity of the soul in accordance with virtue. For, even though this isn't an empty notion, it is clearly a willowy one that becomes fleshed out only by detailed descriptions of how the balanced operations of particular virtues harmonize potentially conflicting dispositions to behavior. Higher-level virtues, denoted by "thin" moral terms, thus supervene on lower-level virtues, denoted by "thick" moral terms.

Such an account, on which higher-level "thin" moral properties effectively supervene on lower-level "thick" moral properties is, I suggest, particularly amenable to elaboration in terms of "homeostatic property clusters," i.e., sets of fundamental underlying causal properties that typically co-occur as a result of a common cause or mutual reinforcement. Boyd hints at this strategy when he advises that we employ property cluster or criterial attribute semantics, which effectively define moral terms by reference to collections of properties some adequate subset of which is sufficient to fix the term's

referent.[6] On such accounts, our concepts of moral kinds are significantly open-textured, exhibiting the essential imprecision or vagueness characteristic of ordinary use. This indeterminacy directly follows from the property cluster form of the relevant definitions, because such definitions are generated not by stipulated necessary and sufficient conditions, but rather by fallible observations of contingent facts regarding explanatorily relevant causal structures. We observe, for instance, that only the proper balance of courage and prudence renders us effective in our pursuit of goals, and we judge that human excellence requires us to self-inculcate proper dispositions to both virtues within a single integrated character.

The operative intuition behind the idea of such a "homeostatic property cluster" account is that various groups of attributes often are contingently clustered as a result of their regular co-occurrence in a statistically significant number of cases, where such co-occurrence is a direct result of these properties' causal homeostasis. Either some of the cluster's properties favor the presence of others under appropriate conditions, or else significant numbers of the clustered properties jointly result from common underlying properties. Clustering may also be causally significant in another way. Important effects may result from these properties only when they are taken in combination, either with each other or with their underlying mechanisms and processes. For example, consider biological species, which clearly illustrate the kind categories that are properly picked out by homeostatic cluster concepts. *Canis lupus*, the species that includes all wolf-like canids including domestic dogs, exemplifies all the features described above. All wolf-like canids share homeostatically related physiological, morphological and behavioral properties. These properties regularly co-occur (e.g., canine teeth and body frames conducive to chasing and hunting). They result largely from common, underlying mechanisms (e.g., domestic dogs differ from gray wolves by, at most, 0.2% of their mitochondrial DNA). And many these properties systematically conspire to place patterned limitations on canine forms of life (e.g., the ability to communicate intention through barking and posture, the propensity for juvenile play and other stereotypical properties that steer dogs into pack social structures). Without overworking the analogy, the moral property account we are sketching here is one in which higher-order moral attributes are importantly similar to biological attributes. "Human excellence" becomes a homeostatic cluster concept given substantive content only by a surrounding cloud of "thicker," intertwined moral properties and relations whose coordinated occurrence constitutes its implementation.

Such an account is recommended by recurrent philosophical intuitions concerning the unity of the virtues. Consider a picture of virtue unity loosely inspired by the work of Nancy Sherman.[7] This picture is particularly commendable for invoking nothing in the way of human essences or final ends. Rather, it starts with the premise there is sufficient commonality among human beings to allow us to articulate approximately true generalizations concerning human desires, generalizations from which we might postulate a representative assortment of apparent goods. Our actual desires, on this account, provide a starting point for our investigations. They launch our inquires in the right direction, as we seek to discern a hierarchy of reliance relationships among our desires in a way that ranks them higher and lower. The items occurring within this range of apparent goods are familiar enough, once again loosely inspired by the work of John Finnis.[8] Identified as non-instrumental, directly pursuable and constituted of activities that are unsusceptible to

assimilation by other activities that would give them greater satisfaction potential than they would otherwise have, these apparent goods consist of the likes of knowledge, friendship, aesthetic activity, practical wisdom and recreation.[9] Presupposing this list for the purposes of the present discussion, the account we have been sketching is one in which we exemplify high-level "thin" moral properties such as human excellence or eudemonistic happiness by substantially participating in acquiring most or all of these lesser goods in an appropriately coordinated fashion. Thus, to live a life of excellence is to exemplify a complex homeostatic cluster property.

Numerous examples support the general outlines of this kind of proposal. Consider, for instance, the mutual reinforcement that must undoubtedly accrue between friendship and the other four items on the short list above. Friendship typically requires actively sharing intellectual, aesthetic or recreational interests, just as the pursuit of these three interests typically requires mutual cooperation, and thus effective sociability with others. Or consider how recreation promotes the kind of physical and mental well-being that is typically necessary for effectively pursuing intellectual and aesthetic agendas. Whatever else may be said of an excellent human life, such a life is likely to be one in which a balance of activities and pursuits allows each component activity and pursuit to be engaged in more effectively than it would be if pursued in isolation. The "thin" moral property of human excellence, happiness, is given substantive content by such "thicker" properties and relations as friendliness and knowledgability, from whose balance and mutual causal sponsorship the property of overall eudemonistic happiness accrues.

On such a cluster-concept account of moral properties, just as higher-level moral properties arise upon on the acquisition of lower-level moral properties that accrue from the satisfaction of various apparent needs, these lower-level properties themselves accrue only from the exercise of many of the traditional virtues. Our ability to acquire primary goods is effectively supported by the likes of courage and prudence, which are necessary for primary goods' pursuit. The virtues, on a standard reading, are habitual strengths of character that promote human flourishing by effectively moderating between states of excess and deficiency in ways that allow us to effectively pursue the primary goods themselves. On the account sketched above, the thesis of the unity of the virtues becomes plausible, particularly when we focus on the dual role that the virtues play in moderating both feeling and action. Think of the questionable effectiveness of courage without perseverance, or without good judgment regarding means and ends, or without a realistic appraisal of our own abilities. This very last example points to the long-recognized fact that, although not all virtues are primary goods, some primary goods are virtues. Consider knowledge (which I have tentatively placed on the list of primary goods), and think of any given range of human activity in which practical wisdom is likely to be useful. For the pursuit of such an activity to be competent, a person must manifest a considered understanding of this activity's value and, therefore, a cognizance of the human benefits provided by this activity. This requires both practical and theoretical knowledge of our own nature, and of human nature more generally, because such information bears upon our

own and others' capacity for pleasure, pain, risk and benefit, as well as the various and sundry tradeoffs that must be made in the course of pursuing and avoiding them. Prudence is made possible only if we possess such knowledge, such as of courage's requirement that risks be taken only for ends that merit them. The same logic goes for kindness and generosity. We can exhibit either of these virtues only if we know enough about people and their circumstances to competently judge what would effectively benefit and encourage them. Thus, knowledge is both an apparent good and a lower-level virtue.

With these stories of how the virtues support our pursuit of a plausible list of primary goods, and of how the acquisition of said primary goods leads to the acquisition of complex relational moral properties, the vocabulary of moral supervenience becomes natural. It becomes tempting to speak of lower- and higher-order moral properties and relations, the higher ones supervening on the lower. This proposal accords with standard conceptions of supervenience on which, for a moral property to supervene on a set of physical properties, it is necessary that the former be present whenever the latter are.[10] This proposal also accords with naturalism, because the former allows particular implementations of moral properties to be simply constituted by occurrent complexes of physical and psychological properties, but without enjoying type-type identity with them. Thus, whatever causal powers a given moral property has, they are inherited from its subvening physical realizers. Moral properties, as such, add nothing new to the causal structure of nature.

With the preceding I have tried to sketch a broad strategy for developing a significant form of moral realism. Using the substantive conceptual content provided by a truth-conditional analysis of our "thick" moral language, we identify a number of human goods that make up a homeostatic property cluster, arrayed above a larger number of subvening goods (i.e., the goods-conducive virtues), themselves literally just identical to baser subvening natural properties, whose implementations literally constitute them, given the situational, attitudinal and perceptual circumstances in which they are realized. I certainly do not pretend that this account is obvious or unproblematic. It presupposes many things; in particular, that some coordinated collections of dispositions are better than others for people to have vis-à-vis the end of human flourishing. It thus presupposes that human flourishing is itself amenable to intelligible and non-arbitrary explication. I have obviously done nothing to demonstrate either claim. Neither have I offered a significant defense of the preceding list of primary goods. However, I think that all three of these omissions are defensible in the context of the present project. A defense of the first two presuppositions would require a book-length treatment, better provided by others. And, once again, the third presupposition is offered only provisionally. Whatever primary goods or goods-conducive virtues we decide upon, all that is really important to the above account is that these goods and virtues be ordered by the sorts of supervenience relations that the account describes. I have not tried to offer a definitive theory. I have merely tried to accommodate some of the seminal tenets of Cornell realism into a story that is substantive enough to allow us to derive some sanguine epistemological morals. But to get these morals we first need to highlight a central feature of moral properties on the preceding account, i.e., their radically relational character.

Relationalism

The preceding section suggests that we identify "apparent" human goods (e.g., knowledge, friendship, aesthetic activity, practical wisdom and recreation) as non-instrumental, straightforwardly pursuable and constituted by activities that are unsusceptible to assimilation by other activities which would give them greater satisfaction potential than they might otherwise have; and that we construe a flourishing human life as one which substantively incorporates at least most or all of these lesser goods. Now, suppose that we also take "intrinsic value" to accrue only to those items that are, in and of themselves, worthy of approval. It clearly follows from these two assumptions that *nothing* counts as intrinsically valuable.[11] This is because moral facts concerning obtaining apparent goods are entirely relative on our account, but only to the extent that they are radically relational.

Consider the candidate excellence-making property cluster alluded to above. Such a cluster is best seen as derived, not through a priori reflection, but through a process of reflective refinement of the actual desires of a particular species of agent, i.e., human beings, and, as such, makes no claims to intrinsic value. Such primary goods, contrarily, are valuable only as they stand in relation to distinctive facts about human nature, where these facts effectively render them relevant by delimiting the range of coherent interpersonal human life plans available for agents to follow. Moreover, relations of mutual reinforcement between various primary goods are themselves likely to be a function of additional relational specifics (e.g., only social creatures with the emotional lives of human beings require friendship relations, fueled only by actively sharing intellectual, aesthetic or recreational interests, the pursuit of which requires, in turn, mutual cooperation, and thus effective sociability with others). Finally, on a subvening level, even a particular virtue, e.g., an agent's courage, occurs, not in that agent's or act's intrinsic nature, but in a relational complex of factors, of which the goal and execution of that act are merely component relata.

The point to note here is that such a construal effectively tames the seemingly fearsome trivializing power of "relativization" talk. Analyzing in relational terms virtues, primary goods and human excellence itself reveals both the misleadingness and ultimate harmlessness of talk about an act's or object's having a certain moral status only from "a human point of view." On such accounts, we can make straightforward sense of how it might be the case that no moral facts could obtain in the absence of human beings performing acts and manifesting states of character to enter relations with the social and natural world around them. But on a relationalist account, such facts do not result from human projection onto the natural order, à la Mackie.[12] Nor are they revealed by a distinctly human quasi-perceptual capacity to discern qualities that pre-exist social reality. Finally, they are not the results of distortive biases through which our pragmatic, conceptual and attitudinal filters color that which would be there in any case. Rather, moral facts require the existence of sentient and reflective human beings, in much the same way that the property of being highly regarded requires the existence of agents and their affections. That an agent's being regarded well by the majority of those who know her requires the existence of other agents and their affections does

nothing to militate against the first-rate realistic status of the fact that the agent is, in fact, highly regarded. It does nothing to show that this isn't one more fact about this agent that we are obliged to include in any objectively accurate description of all there is to be said about her. Similarly, that an act can be praiseworthy for its courageousness only when manifested by such reflective and vulnerable creatures as human beings, as opposed to, say, angels or spirits, does nothing to show that the act's praiseworthiness is a fabrication of the human mind or heart. Rather, it merely highlights the fact that no relation can obtain in the absence of its proper relata.

We might thus suggest that this relational status of moral properties should hardly come as a surprise to us. It is, in fact, bluntly recommended by the radical multiple realizability that is characteristic of moral supervenience. Granted, we sometimes do, when talking about candidate virtues and intrinsic goods, speak of moral values as though they were monadic properties. However, this is best viewed as a convenient shorthand. Relational properties are often treated as monadic properties only because it is convenient to view them as relativized to standard or normal background conditions that we are inclined to treat as given (e.g., we say that seven-foot George is tall only because we take it as a given that the folks to which he is being compared are generally shorter than this). Similarly, the goodness of apparent goods is recognizable only insofar as those goods function as they do within the relational matrices of human needs, desires and capacities that they occupy. We need not suppose, on this account, that anything's value must shine from within, as it were, like Kant's jewel of right intention.

When we so recognize the relationalist character of moral facts and properties, we answer many questions. One of these is the *access question*: How can we perceive the occurrence of occurrent moral properties? On the account above, the answer is that we can't. Or, at least we can't perceive occurrent moral properties any more or less than we can perceive other occurrent relational properties, such as that of "being older than." Relations generally should be viewed as ontologically dependent upon the monadic properties of their relata. Thus, we can perceive that George is older than Michael without registering the occurrence of a monadic or intrinsic quality of "being older than" which is somehow inherent in George. George's occurrent property of being older than Michael is an occurrent relational attribute that functionally supervenes on the monadic and intrinsic qualities that George and Michael have, of being 41 and 40 years old, respectively. Similarly, I suggest, Oscar Schindler's courage consists of his complex of dispositions to do such things as lie to the Gestapo, but only as these dispositions stand in relation to such items as his intentions, sobriety and situation of genuine risk. We thus might want to say, à la Grundlagen-vintage Frege, that courage, though "objective" (i.e., properly amenable to being taken as a subject by true or false sentences), is not "actual" or "concrete" (i.e., properly detectable by the senses). But in these respects, courage is no different from numerous other clearly respectable properties and relations, none of which can be sensorily detected, after all. Rather, their implementations can be sensorily detected, but only when we have sufficiently widened our field of view. There is, on the face of it, no reason to think that the facts that make for an implementation of courage are any less actual than the facts that make for implementations of other, more clearly natural, properties and relations.

To the *location question* of where moral attributes are, the relationalist account answers, very simply, "All around us." Again, this response is made plausible by the fact that moral attributes, far from being stationed in some mysterious non-causal realm, are relational composites of mundane natural attributes, whose normative import is recognizable not through occult powers of "goodness," but through the roles they play in allowing virtues to be exercised, primary goods to be acquired and higher level excellence-constituting homeostatic property clusters to be implemented. Again, envision Schindler lying to the Gestapo. The courageousness of this action consists not of some unnatural superlunary property of goodness, but in a complex of natural constitutive features which collectively function to constitute a homoeostatic property cluster, in this case a stable character with which Schindler lives.

Correlatively, this relationalist line of response to the location problem does much to answer a long-standing concern over the tenability of moral supervenience. This concern is with the following question. What is the relation of moral supervenience supposed to come to, given that it is conceded to be clearly distinct from the relation of logical entailment? A standard answer to this question is offered by Platts. Subvening facts "fix" supervening facts, where "fixing" is explained largely through analogy with other, more prosaic, non-moral part/whole relations. There is a problem with this strategy, however, as Platts himself admits. It leaves him with few analogues with the fixing relation to which we can appeal in our efforts to render moral emergence intelligible and familiar.

Consider Platts' example of pointillesque portraits.[13] Envision a two-dimensional figure of a human head, comprised of clumps of black dots on a canvas. Here, Platts initially suggests, we have a clear example of the "fixing" relation. If we alter the pointillesque distribution, we may vary the head's shape. Similarly, if we alter the head portrait's shape, we automatically change the pointillesque distribution. And yet, even though *some* spatial description of the distribution of pointillesque components may, in fact, imply that the painting is that of a head, it is not *through* this description that we perceive the painting as a head portrait. Even though the dots fix the figure, we notice the latter independently of the former.

Platt originally hoped to use this example as a precedent for the kind of non-inferential supervenience relation with which he aspires to elucidate moral properties. On such an account, the relationship between Oscar Schindler's courage and its subvening circumstances might be one in which the former, even though consisting of nothing more than Schindler's disposition to falsify records in pertinent circumstances of risk and sobriety, may not be inferred from these subvening factors. Unfortunately, Platts admits, this hope is forlorn, for the cases are deeply disanalogous. In the Schindler case, it is clearly necessary to validate judgments of the man's bravery by citing the aforementioned facts concerning his dispositions and situation. These facts comprise the very reasons that Schindler's act is courageous. In this, Schindler's case is unlike that of Platts' pointillesque portrait, where it is unnecessary to validate judgments about the figure's shape by citing facts concerning constitutive dots' spatial distribution. The worry that emerges from this difference is that obtaining the fixing relation between, e.g., dots and pointillesque portraits provides no precedent for the alleged supervenience relations that arise between, e.g., Schindler's dispositions and circumstances,

on one hand, and his courage on the other. Because the former example is, as the latter is not, one in which allegedly supervenient properties are recognizable without reference to their allegedly subvening bases, only in the former case does the fixing relation seem to give us what we want, i.e., a way for a supervening property to enjoy a claim to autonomous reality, a life of its own, as it were, beyond those of its constitutive subvening basal properties.[14]

Fortunately, however, a relationalist construal of moral propertyhood allows us to acknowledge this difference between the two examples without deriving any conclusions from it that support moral property attribution's alleged second-rate status. For, on a relationalist account, it is easy to explain, in principle, why our ascription of courage to Schindler obliges us to itemize the facts which constitute Schindler's implementation of courage. The relational features of occurrent circumstances that make for this implementation simply cannot be taken in at a glance in the way that the relative placement of black dots making for the occurrence of a pointillesque portrait can be. We need to justify the occurrence of courage by citing relevant subvening facts because these facts extend into the past and the future, and include non-straightforwardly observable details concerning personal history, intended aims, motives, circumstances of risk, states of mind and self-conception. Consequently, we have no reason to interpret this feature of moral facts (i.e., their reliance upon determining features that are difficult to observe all at once, due to their mentalistic nature and/or dispersion across space and time) to indicate their unreality or lack of autonomy. Why shouldn't we instead take this obvious feature of moral facts to illuminate the above-cited suggestion that moral facts are relational composites of mundane natural attributes, whose normative imports accrue to them only by virtue of their causal positions within a relational matrix of primary goods and supervening property clusters of potentially dizzying complexity?

Externalism

This sort of relationalist analysis can be cautiously recast in externalist terms. The moral meanings of actions are likely, though not guaranteed, to be externalistically delineated to roughly the same extent that they are relationally constituted. For, if a fact consists of selected patterns of relata which are themselves largely external to an agent, the more likely it seems that such facts would obtain outside the agent's conscious awareness.

Here we must walk a fine line, however. Moral facts, as such, enjoy normative, rather than merely causal, efficacy. And normative efficacy involves following rules, not merely obtaining laws. But following rules, by its nature, cannot be completely unwitting. Otherwise, to use Sellars' oft-quoted example, thunder could be said to track lightning as a result of following a rule.[15] Thus, we might suspect that tension exists within the above account between its relationalist and normative aspects, about which it has said nothing. It is time to address this concern. We can do this below in two stages. First, we can argue that moral

discourse is especially amenable to externalistic construal. Second, we can argue that the anticipated tension does not arise once we segregate issues of moral worth from issues of moral responsibility.

As for the first point, the argument is comparative. Moral discourse is especially amenable to externalistic analysis as compared to, e.g., psychological discourse. Beyond concerns about dead and gone historical facts' capacity to exercise casual efficacy in the present (except insofar as they have been somehow intrinsically encoded into occurrent mental and semantic events' internal natures), the most serious reservations about externalist semantic analysis are expressed by philosophers who note that we have no good reason to invoke broad content because we have no good reason to suppose, e.g., that the thoughts expressed by the Doppelgangers of familiar Twin Earth thought experiments differ from each other in any way at all.[16] These characters' apparent behavioral dispositions are, after all, exactly the same. And even though one is led by his verbal dispositions to H_2O and the other is led by his verbal dispositions to XYZ, this has nothing to do with content and everything to do with contingent contextual circumstance. We may state this point alternatively in evolutionary terms by elaborating on Fodor's observation of how "empirical theories are responsible only to generalizations that hold in nomolologically (as opposed to conceptually) possible worlds."[17] Given that Jones grows up on Earth and Twin-Jones grows up on Twin Earth, we should immediately wonder what evolutionary mechanism could ever give either character an ability to determinately mean, or refer to, one substance as opposed to the other. What evolutionary pressures could favor a specific over a vague or equivocal referential capacity in such cases? It is clear to philosophers, of course, that there should be differences in content, but this is only because philosophers consciously stipulate a descriptive distinction between the two worlds and substances that gives each term a distinct and narrow content. The conclusion that broad content must be operative is never, in cases of such thought experiments, the result of anything more than an implicit assumption that the characters in the externalists' scenario must, by virtue of externalist content, make the very same distinctions that are available to philosophers by virtue of their own internalist content distinctions.[18] An equivocation of perspective and subsequent "informational mismatch" is at play in such thought experiments, because they encourage us first to assume a God's-eye point of view on the subject's referential situation, and then to attribute, on the basis of this unearned point of view, semantic properties to him which, properly speaking, accrue only to us.[19] This is not to suggest that a perfectly useful function isn't served by likening speakers to texts, about whom interesting questions can then be posed concerning the "deeper meanings" that might result from social and natural factors of which the speaker is unaware. The claim here is merely that such externalist ascriptions of meaning can be, at most, supplements to internalist ascriptions. They cannot effectively replace these internalist ascriptions for the purposes of providing a fully adequate account of the varieties of legitimate meaning.

On this fact a marked difference separates the property concepts of semantics from those of ethics. Moral notions are much more likely candidates for thoroughgoing externalist analysis than are semantic notions. Even when we consider the role of "right intention" in moral action, we need

not find it to take the form of conscious access to an introspectable intention to "do the right thing." Once again, consider the strange case of Oscar Schindler. In the account we have been painting, Schindler's disposition to falsify records effectively fixes his manifestations of courage only given such internal mental states as the realistic sobriety of mind and the personal perception of the genuinely risky circumstance in which he manifests it. But these mental states are not the point. For Schindler's actions to be courageous, it is not required that he perceive them as such. And this, I take it, is what would be mandated by an internalist analysis of courageous action, an analysis in which an agent's courageousness requires that he consciously know or intend his action to be courageous (just as, in this account's sister semantic internalist construal, an agent's denoting dogs by the word "dog" requires that he consciously intend to pick out such and such kind of animal through the mediation of such and such descriptive content). Intending to do the "right thing" is not a prerequisite for moral action. Think of the virtues Huck Finn displays when he decides that he simply must help runaway Jim *despite* his conscious conviction that this action is immoral (let alone his literal expectation that he can expect nothing but an eternity of righteous hell-fire in return).

Thus, whereas there are *indispensable* internalistic construals of what we mean when we talk of dogs or water, there are no corresponding indispensable internalistic construals of what we mean when we talk about what we need, what we should pursue, how we should be or what we should do. Rather, there are essentially internalist notions of moral responsibility which make it the case, e.g., that an agent cannot be held responsible for actions he performs due to non-willful ignorance and the like. Importantly, the distinctions at issue here, between our moral assessments of acts and characters on one hand, and of agent culpability on the other, are long-standing ones. We commonly allow, for example, that a 12-year-old Khmer Rouge revolutionary may have done unspeakable things for years and may have lacked anything remotely resembling a coherent set of excellence-making personal virtues while he did them, yet we may not blame him for his youthful and ignorant crimes. What this illustrates is that the essentially externalist bases of our judgments of virtue and obligation are clearly separable from the essentially internalist basis of our judgments of moral responsibility and blameworthiness. Because this operational distinction already exists in our moral reasoning, there is no danger that an externalistic analysis of other evaluative terms will cause us to ignore indispensable internalistic aspects of our moral reasoning. But in this the moral case is crucially different from the semantic one, where confusion seems to constantly arise precisely because the proper functions of internalistic and externalistic analyses are not clearly segregated while being simultaneously, if only implicitly, acknowledged.

Moral Knowledge

The preceding two sections sketch selective aspects of an account on which moral property terms have externalistically determined normative content, and moral properties are relational properties whose normative import is incumbent upon their causal positions within a relational matrix of primary goods and

supervening property clusters of extreme convolution. Once again, however, this paper's reasons for sketching metaphysics are ultimately epistemological. Moral knowledge seems more problematic than prosaic physical knowledge, and any adequate account of moral factuality should explain why. The above account helps to provide such an explanation.

Recall our comparative remarks concerning Mark Platts' pointillesque head portraits and Oscar Schindler's courage. As we noted, there is, contra Platts, a very simple reason that we can simply perceive the representational aspects of a portrait, but not the moral aspects of an action; and that judgments of the latter sort, but not the former, require us to explicitly recite the reasons that justify our verdict. The reason, again, is this. The relational features of occurrent circumstances that make for an implementation of courage simply cannot be taken in at a glance, as it were; whereas the relational spatial features that make a pattern of dots a portrait of Winston Churchill can. The epistemological moral of this is probably obvious. Discerning the occurrence of a moral fact is never as simple as perceptually registering the occurrence of a prosaic physical event. Even though the sensory faculties through which we detect the features of facts are precisely the same in both cases, the pattern-recognition mechanisms through which we selectively collocate these features differ in important ways.

To appreciate how and why they differ, we need to turn from the similarities between higher-order moral attributes and biological kind attributes emphasized so far, to the obvious dissimilarities that we have downplayed or ignored. For, to the extent that the class of persons is not a merely biological category, such differences are to be expected. Human lives, like other biological lives, take their organic unity from the homeostatic property clusters that mold their defining processes of growth and development into organic wholes over time. But it is philosophically commonplace that human beings, as opposed to most or all other creatures, are extra-biological in that their organic unity is determined by psychological, as well as physical, factors.[20] And it is hard to imagine that this aspect of their continuity could be captured except by recourse to the personal narratives through which peoples' past and present actions and future goals are relationally located as components of coherent wholes. To recognize moral facts *as such* nothing need be added to the pertinent perceptual data. What we must contribute, rather, are narrative tellings and retellings of events which serve to bring their inter-relational normative significance to light.

It is for this reason that the recognition of a moral fact, supervening or subvening, is never as simple as perceptually registering the occurrence of a prosaic physical event. Neither is it as simple as discerning the occurrence of a causal connection. To recognize a moral fact is to describe a relational narrative to ourselves. We must note that "Schindler's mouth moved in such and such a way, as he stood face to face with such and such a person, each of whom had such and such personal and causal histories and self-conceptions, all within such and such an institutional setting . . . and that's why Schindler's lying to the Gestapo officer was courageous." By way of literary example, think of Nussbaum's wonderful defense of Odysseus' decision to forego immortality and transcendence in the arms of Calypso.[21] Odysseus, she notes, suffers from no illusions. He is fully aware that the Goddess surpasses Penelope in form and stature and that the quest for transcendence is not without its own moral appeal. But despite this,

> [Odysseus] opts to continue his voyage, thus choosing not only risk and difficulty, but the certainty of death. And not only death, but the virtual certainty that he will at some time lose what he most deeply loves, or will cause by his own death great grief to her. He is choosing the whole human package, mortal life, dangerous voyage, imperfect mortal aging woman. He is choosing quite simply what is his, his own history.[22]

We need not presume that there is anything of intrinsic value in Odysess' decision to forego immortality. Rather, we would do better to suppose that only in the context of Odysseus' narrative recounting of his own life does this decision become an instance of right action. The fact that his love-making with Penelope is a prelude to conversation, whereas his love-making with Calypso seems to provide a mere segue to sleep, nicely reflects this fact, because the shared narrative of their lives gives the former pair something to talk about.[23]

Or, more gloomily, think of Winston and Julia in the final pages of *1984.*

> 'Sometimes,' she said, 'they threaten you with something — something you can't stand up to, can't even think about. And then you say, "Don't do it to me, do it to somebody else, do it to So-and-so." And perhaps you might pretend, afterwards, that it was only a trick and that you just said it to make them stop and didn't really mean it. But that isn't true. At the time when it happens you do mean it. You think there's no other way of saving yourself, and you're quite ready to save yourself that way. You want it to happen to the other person. You don't give a damn what they suffer. All you care about is yourself.'
>
> 'All you care about is yourself,' he echoed.
>
> 'And after that, you don't feel the same toward the other person any longer.'
>
> 'No,' he said, 'you don't feel the same.'
>
> There did not seem to be anything more to say. The wind plastered their thin overalls against their bodies. Almost at once it became embarrassing to sit there in silence: besides, it was too cold to keep still. She said something about catching her Tube and stood up to go'.[24]

In this passage there is, subtly but clearly, amid the darkness of utter despair, a glimmer of hope in the mere fact that both characters retain sufficient humanity to feel embarrassment at their situation. Though a small victory in the face of crushing spiritual defeat, it speaks well for Winston and Julia that they can still experience discomfiture at the enormity of their loss, out of mutual disappointment at their own frailty and failings before their own ideals and each other's. Moral value is implemented in the prosaic fact that the two uncomfortably share a bench momentarily with downcast eyes before standing to walk away. In the course of this sequence of physical and psychological facts they manifest their sharing of a morally appropriate sentiment that not even the torturers could take from them.

The point at issue concerns the manner in which narrative context renders the evaluative status of prosaic events detectable. It does this, once again, not by adding to the content of such perceptions or by projecting an arational Humean emotional covering veil of approval or disapproval over them. Rather, it

highlights various relations between prosaic facts so as to render these collective facts digestible to moral perception. The worthiness of Odysses' decision to forsake immortality is not available to unstructured perception. Neither is that of Winston's and Julia's tense embarrassment. It is available only to perception that has been sensitized by precisely the kind of narrative that Homer's and Orwell's novels provide.

It is easy to imagine this kind of account falling on skeptical ears. For isn't it largely *because* we can recognize moral facts only within the context of complicated relational narratives that we are inclined to doubt their objective status? A moral skeptic's suspicion is evident in the following: "Aren't we *just* telling stories with such narratives, after all, and shouldn't we be wary of our propensities for story-telling?" To such skeptics, however, we should insist that the above relationalist account aims to circumvent such doubts by arguing that the moral skeptic's demands were always inappropriate to their particular subject matter. For what the relationalist account reminds us of is that we never should have expected moral facts to be straightforwardly observable. Such facts are recognizable, and hence discoverable, only through the telling of stories that human beings find resonant, because only through such narrative accounts can these facts' component relata be assembled into an intelligible form.

Think of Hume's suggested indicators of occurrent causal relations in the *Treatise*: contiguity, priority and constant conjunction. Causal connections, he effectively tells us, are things whose component determining elements are, by and large, readily encompassed under a single act of inspection. We attribute such connections to event pair types whose instances come to us mostly pre-assembled, lying next to each other in space and time. The courageousness of an agent's action, on the other hand, is incumbent upon relata concerning history, intentions, states of mind, circumstances of risk and the like, which generally do not so readily converge for easy sensory intake. Or consider Hume's own familiar example. "There is no fact of the matter, or real existence," he insists, that can be perceived in an act of "willful murder" that identifies it as "vicious." Hume's mistake is to draw skeptical conclusions from examples such as this. We should not expect such an occurrence, taken in isolation, to wear its moral status on its face.

Nor should we endorse Hume's programmatic response to these failed expectations, because he reasons in the following fashion. Because we cannot discern evaluative qualities within the act, we should look instead within ourselves. Thus he concludes in the Treatise, "In whichever way you take it, you find only certain passions, motives, volitions and thoughts. . . . The vice entirely escapes you, as long as you consider the object . . . You never can find it, till you turn your reflection into your own breast . . . ."[25] It is thus that emotivism or simple sentimentalism becomes such a tempting option for those who first seek to discern moral qualities within particular actions, considered in analytic isolation. The point at issue here is not merely that emotivism is anathema to the kind of cognitivist position we have been examining. Rather, it is that an even deeper flaw mars Hume's prompting strategy of reflecting "into his own breast." For if the preceding relational analysis is correct, semantic externalism is particularly appropriate to the analysis of moral language, rendering hopeless any such purely introspective inquiry.

Conclusions

  In short, the features that make Oscar Schindler's actions morally resonant in the first place can be gleaned only from the sort of narrative account that skeptics are inclined to dismiss as "mere story-telling." To the extent that moral skepticism stems from a Humean conception of what moral factuality would have to be like if it were to obtain, moral skepticism is misplaced. To look for virtue or vice by closely examining the exercise or products of human agency, considered in isolation, is akin to scanning for giraffes through a microscope. It is a made-to-order failure of a kind with the classic Indian fable of the six blind men examining the elephant. What story-telling, or narrative, accomplishes for knowledge of moral factuality is, we might say, a type of fine-tuning of the examinative focal length, a process no less necessary for discerning goodness than it is for spotting planets.

Notes

[1] Brink, David (1984) "Moral Realism and Skeptical Arguments from Disagreement and Queerness," *Australasian Journal of Philosophy* 62; Brink, David (1989) *Moral Realism and the Foundations of Ethics*. (Cambridge: Cambridge University Press); Boyd, Richard (1988) "How to Be a Moral Realist" in *Essays in Moral Realism*. ed., Sayre-McCord. (Ithaca, NY: Cornell University Press); Railton, Peter (2003) *Facts, Values, and Norms: Essays toward a Morality of Consequence*. (Cambridge: Cambridge University Press); Sturgeon, Nicholas (1985) "Moral Explanations" in *Morality, Reason and Truth*, ed., David Copp and David Zimmerman, (Ithaca: Cornell University Press); M. Platts (1979) *Ways of Meaning* (Cambridge, Mass: MIT Press).

[2] Davidson, Donald (1990) "Radical Interpretation" in *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press); 'The Structure and Content of Truth' (1989), *Journal of Philosophy* 87.

[3] This is all understood under the qualification of semantic holism, of course. What is ultimately primary in such accounts is overall theory, against which T-sentential and referential ascriptions are only relatively settled.

[4] McDowell, John (1980) "On the Sense and Reference of a Proper Name" in *Reference, Truth, and Reality* (London: Routledge and Kegan Paul); Platts, Mark (1979) *Ways of Meaning*, (Cambridge, Mass: MIT Press), Ch. 10.

[5] Platts, Mark (1979) *Ways of Meaning*, pp. 225, 246.The locution of "thick" and "thin" moral concepts was originally articulated by such arch moral antirealist as Bernard Williams in Williams, Bernard (1985) *Ethics and the Limits of Philosophy* (Cambridge :Harvard University Press), pp. 144-145.

[6] Boyd, Richard (1988) "How to Be a Moral Realist," in *Essays on Moral Realism*.

[7] Duncan, Steven M. (1995) *A Primer of Modern Virtue Ethics* (Lanham, Md.: University Press of America); Sherman, Nancy (1989) *The Fabric of Character: Aristotle's Theory of Virtue* (Oxford: Clarendon Press).

[8] Finnis, John (1980) *Natural Law and Natural Rights* (Oxford: Oxford University Press).

[9] Duncan, Steven M. (1995) *A Primer of Modern Virtue Ethics*.

[10] Kim, Jaegwon (1998) *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation* (Cambridge, Mass.: MIT Press), p. 9.

[11] A contrasting account of intrinsic value, which many have endorsed in broad outline, is particularly well developed in Lemos, Noah (1994) *Intrinsic Value: Concept and Warrant* (Cambridge: Cambridge University Press), p. 15.

[12] Mackie, J.L. (1977) *Ethics: Inventing Right and Wrong* (Harmondsworth [England]; New York: Penguin).

[13] Platts, Mark (1979) *Ways of Meaning*, pp. 244, 251, 253.

[14] See Arrington's articulation of this difficulty, which is somewhat clearer than Platts' own. Arrington, R. L. (1989) *Rationalism, Realism, and Relativism* (Ithaca and London: Cornell University Press), p. 147.

[15] Sellars, Wilfred (2000) "Does Empirical Knowledge Have a Foundation?" in *Epistemology: An Anthology*, ed. Ernest Sosa and Jaegwon Kim (Oxford: Oxford University Press), pp. 120-124.

[16] Martin, C.B. and Heil, J. (1998) "Rules and Powers," *Philosophical Perspectives 12*, pp. 383-312.

[17] Fodor, Jerry (1994) *The Elm and the Expert* (MA: MIT Press). Needless to say, Fodor may not agree with this use of him in the context of an evolutionary argument.

[18] As in Davidson's swampman case, where the swampman's thoughts and words must refer to nothing, we are told, because it is obvious that "what a person's words mean depends in the most basic cases on the kinds of objects and events that have caused the person to hold the words to be applicable; similarly for what the person's thoughts are about," Davidson, Donald (1987) , "Knowing One's Own Mind," *Proceedings and Addresses of the American Philosophical Association 60*, pp. 441-58.

[19] The locution here (of "informational mismatch") is due to Fogelin, who uses it in a very different, epistemological context in Fogelin, Robert (1993) *Pyrrhonian Reflections on Knowledge and Justification* (Oxford: Oxford University Press).

[20] This is, of course, not to suppose that psychological properties and facts differ in kind from physical properties and facts. It is to suppose only that the former constitute a distinct variety of the latter.

[21] Nussbaum, Martha (1992) *Love's Knowledge: Essays on Philosophy and Literature* (Oxford: Oxford University Press), pp. 365-367.

[22] Ibid., p. 366.

[23] Ibid., p. 367.

[24] Orwell, George (1949) *Nineteen-Eighty-Four*. Retrieved August 20th, 2006 from Project Gutenberg of Australia online edition at http://gutenberg.net.au/ebooks01/0100021.txt.

[25] Book III, Part I, Section I.