

THE RULES OF “GOODNESS”: AN ESSAY ON MORAL SEMANTICS

Paul Bloomfield

Here is a story about the word “good” that has been promulgated by empirically minded metaethicists living at one distance or another from Ithaca, New York.¹ These “Cornell realists,” for want of a better name, use the causal theory of reference to explain how the word “good” gained its reference, or got its meaning, co-opting for metaethical purposes the persuasive moves that Putnam and Kripke used to explain how “water” got its meaning.² The briefest summary of the familiar story is that there was an intention of a linguistic community to refer to water, the physical substance that we later learned is H₂O, with the word “water.” Thus, it was H₂O that causally regulated the use of the word “water,” even before we discovered that water is H₂O. When we did discover water’s identity to be H₂O, we could then be said to have gained knowledge that is necessary *a posteriori*. If the Cornell realists are right, then there is a property or group of properties that causally regulate the use of the “good” (or “right”), even if we, as empirical investigators, have not yet discovered which properties these actually are. When we do eventually make this discovery, presumably ending the debate that constitutes normative ethical theorizing, we will find

that goodness is necessarily identical to those properties, just as water is necessarily identical to H₂O. And as our knowledge that water is H₂O is necessary *a posteriori*, our knowledge of what good is will thereby merit the same secure status.

If this is the case, then our discourse should have a certain semantic commitments about what we are willing to say about things like water and goodness. Unfortunately for the Cornell realists, it seems as if our intuitions about the semantics of “water” match up with what we would expect, if water were (necessarily) identical to H₂O, but the same cannot be said about the semantics of what is morally “good.” The semantics which emerges from the moral ontology of the Cornell realists is not the semantics that we ordinarily employ in our moral discourse. If this problem is generalizable to all empirically driven versions of moral realism, then those realists are in grave trouble. If, however, there is a version that avoids this semantic difficulty, then that alone is much to speak of in its favor.

The “new wave” semantic problem has been made precise in an extended series of papers co-written by Terry Horgan and Mark Timmons.³ In order to explain it, the

place to start is with Horgan and Timmons' analysis of the Cornell realists' semantic commitments. This analysis is a reconstruction of the semantics provided in Boyd's version of moral realism (op. cit.), which is intended to capture the essentials needed for generating the problem (though, as described below, Horgan and Timmons have shown it to be impressively generalizable). Boyd (1) construes terms like "good" as natural kind terms, (2) having synthetic definitions yielding the essence of the property "goodness" (presumably this would something like a Lockean "real essence"), and these definitions would (3) have the status of necessary *a posteriori* truths. If we take "water" as our model natural kind term, then by employing a Twin Earth thought experiment (also described below) we will learn that it rigidly designates H₂O, and hence (4) "good" will also rigidly designate the property or group of properties to which it refers. This reference is gained by causal connections formulated by Boyd as follows:

Roughly, and for nondegenerate cases, a term *t* refers to a kind (property, relation, etc.) *k* just in case there exist causal mechanisms whose tendency is to bring it about, over time, that what is predicated of the term *t* will be approximately true of *k* (excuse the blurring of the use-mention distinction). Such mechanisms will typically include the existence of procedures which are approximately accurate for recognizing members or instances of *k* (at least for easy cases) and which relevantly govern the use of *t*, the social transmission of certain relevantly approximately true beliefs regarding *k*, formulated as claims about *t* (again excuse the slight to the use-mention distinction), a pattern of deference to experts on *k* with respect to the use of *t*, etc....When relations of this sort obtain, we may think of the properties of *k* as *regulating* the use of *t* (via such causal relations). (p.195)

We will discover which properties regulate "good" by pursuing debates in normative ethical theory, and presumably these debates will be ended by discovering which properties are doing this work. Thus, Horgan and Timmons take the Cornell realists to be committed to the formula they have baptized the "causal regulation thesis" which, when combined with (1)–(4) above, generates "causal semantic naturalism":

CRT, Causal Regulation Thesis: For each moral term *t* (e.g., "good"), there is a natural property *N* such that *N* alone, and no other property, causally regulates the use of *t* by human beings.

CSN, Causal Semantic Naturalism: Each moral term *t* rigidly designates the natural property *N* that uniquely causally regulates the use of *t* by human beings.⁴

With these commitments on the table, we are now asked to compare a visit to a Twin Earth filled with XYZ and not H₂O with a visit to what Horgan and Timmons call "Moral Twin Earth." If we encounter a Twin Earth with stuff that very much looks like water but is actually XYZ and not H₂O, most people (so it seems) would think that the word "water" has two different meanings on Earth and Twin Earth, because the word is causally regulated by two different kinds of stuff.⁵ Our linguistic intuitions tell us that the meaning of "water" is rigidly designated by H₂O, where that means that even before we knew what water was, we took "water" to refer to *that* substance at any possible world where it exists and to lack reference at all other possible worlds. Now, consider the analogous case with "good." Suppose that here on Earth we end the debate in normative ethical theory with the discovery that the properties actually mediating our use of the term "good" are those that traditional consequentialists have maintained. Goodness then would be necessarily identical to

(roughly) the maximization of happiness and "good" would rigidly designate what instantiates these properties. We then proceed to Moral Twin Earth, and, lo!, find that these Twin Earthlings have discovered that the term "good" rigidly designates, and is causally regulated by the properties that deontologists have maintained, and goodness (here?) is necessarily identical to what instantiates the property of (roughly) acting from the motive of duty. Now, it seems unjustified and hegemonic (imperialistic?) to say that the Twin Earthlings have got it wrong in regards to the meaning of "good," while we here on Earth got it right; yet it cannot be the case that goodness is necessarily identical to what instantiates both the maximization of happiness and acting from the motive duty. Nor can it be the case that "good" rigidly designates two properties; this is to claim that "good" is not like "jade" which refers to both jadeite and nephrite.

Our intuitions (moral or semantic?) disallow a reconciliation between us and them by saying "Oh you can use 'good' as you do and we'll use it as we do, and 'good' will have different meanings for us," which is how our intuitions tell us to behave with regards to the meaning of the word "water": assuming we are all of the same biological species and English is identical in all other respects, we would consider it as a trivial homonym if Twin Earthlings call "water" something that is not, but we could not settle for that diagnosis if they call "good" something that is not. Thus, the Cornell realism fails, for at its heart is the CSN account of the meaning of moral terms, and this account fails: ordinary moral discourse would say that we have (at the very least) a semantic disagreement with the Twin Earthlings about the meaning of "good" and we would not say the same about "water." If "good" were like "water," then we should not find these differences between them. Any form of moral

realism that employs a semantics of moral terms putatively modeled on those of "water" fails (to that degree) to comport with fundamental intuitions about how moral terms may be used. Horgan and Timmons are right: any version of moral realism that comes to this result is doomed.

One might wonder, however, why anyone would be surprised at the failure of modeling a semantics of "good" on "water." Ontologically, there is a great difference between goodness (as a property) and water (as a substance); there seems little *prima facie* reason to think that semantic analyses of the uses of "good" and "water" would reveal two analyses which have the same structure. There are many ways in which the semantics of "good" and "water" obviously differ. Just consider how each handles negation, how the shape of the extension of "not good" is vastly different than that of "not water"; being not good comes in degrees while being not water does not. "Good" has an antonym, "water" does not. Semantic difficulties like these surface all over. Discourse about "good" is far more complicated and manifold than discourse about "water."

At least some of these complexities can be captured by noting that the semantics for "good" has at least three different uses of the term. To work up to these three uses, begin with the fact that everyday moral discourse contains predications of "goodness" both to people (or lives) and to acts; in ethical theory, this distinction is manifested by the debate between "agent-centered" and "act-centered" moralities.⁶ Agent-centered morality focuses on the predications of "good" to people or lives, and speech acts in which we say "Mother Teresa was a good person" or "Socrates lived a good life" are ordinary enough. Typically, focus on this use of "good" leads one to study virtue ethics, as opposed to either consequentialism or deontology.

Regarding act-centered predications of “good,” we can distinguish two types by noting that they take different prepositions. On the one hand, we commonly say that certain acts are good *for* people, while on the other we say of certain acts that it is good *of* people to do them, or that they are done *from* a sense of fair play, or *from* the motive of duty. When the preposition “for” is attached to a predication of “good,” what we are saying is that the act is being called good in virtue of its instrumentality, its causal effects or consequences, or what the act is good for accomplishing, or obtaining, or doing. We say that “doing charity work is good for a person,” insofar as it may causally contribute to that person’s well lived life, as well as contributing to the well lived life of the receiver of the charity. Of other acts, we call them “good” because of what they signify about the agent performing the act (or perhaps the relevant motives *of* the agent). In such cases, we use often use either the preposition “of” or “from,” both of which are indicative of the etiological source of the act (as opposed to the consequences of the act). Thus, we can say to a friend paying an unexpected and kind visit, “it is so good of you to come,” indicating that the visit is some sort of proof (or a sign) of the person’s good will toward us. John Rawls is on this track when he talks about acts done “from a sense of fair play”; Barbara Herman’s paper title uses this sort of predication, “On The Value Of Acting From The Motive Of Duty.”⁷

It should be no surprise to anyone by this point that these different predications of “good” are the verbal manifestations of two different ways of thinking morally and evaluating acts; it is the semantic embodiment of the distinction between consequentialist and deontological moral thought. It should not be surprising that such a distinction exists in everyday moral language,

nor should it be surprising that “good” can be used in both ways; in everyday morality, we sometimes evaluate acts *qua* consequentialists and other times we do so *qua* deontologists. The distinction just outlined between predications of “good” on acts, marked by different prepositions, is itself good evidence that moral thought is complicated in just this way. For this reason alone, we should not be surprised that the debate between these normative ethical theories has proved to be so recalcitrant: these uses of “good” do not seem to be reducible to a single use. Sometimes we argue over the morality of acts like consequentialists, other times we argue like deontologists, and still others we argue over the moral qualities of agents, like virtue theorists. What we need is an account of moral terms, and an ontology to hang the account on, that can support talking like consequentialists, deontologists, and virtue theorists. One ought to expect that the semantics of moral terms is going to be a more complicated business than the semantics of substance terms, like “water” or “gold.”

The Moral Twin Earth problems were formulated with the Cornell realists in mind, but it turns out that other versions of moral realism have analogous problems. Horgan and Timmons call this a “recipe” for a problem for moral realism.⁸ Michael Smith’s “metaethical rationalism” has similar troubles, as do the positions of David Copp and Frank Jackson.⁹ And there is a version of the problem, for those holding that values and colors have the same ontological status, namely as secondary qualities equally fit for a dispositionalist’s treatment.¹⁰ What moral realism needs is an ontology that does not stimulate such semantic headaches.

And the place to begin in looking for a solution to the problem is to note an underlying assumption made by the irrealists who pose the problem: this is to suppose

that realism and relativism are somehow at odds or are incompatible with one another. Indeed, one need only look at the titles of Horgan and Timmons' paper "From Moral Realism to Moral Relativism in One Easy Step" (as if the slide to relativism is automatically problematic) and Mackie's argument, the "Argument from Relativity."¹¹ This unhappy bit of metaphysics is surprisingly common. Consider how our best proponent of modern moral relativism, Gilbert Harman, lays out the metaphysical territory.¹² He considers moral relativism to be the view that morality "arises from some sort of convention or understanding among people, that different people arrive at different understandings, and that there are no basic moral requirements that apply to everyone" (p. 27). Thus, he groups moral relativism among the other moral non-realisms (non-cognitivism and skepticism), and these are contrasted with absolutism and (presumably, thus) realism. Harman, taking "naturalism" to be the rightful heir to empiricism, thinks naturalism leads to relativism et al., while absolutism is the errant result of treating ethics as "autonomous." Boyd and other Cornell Realists, even those who call themselves "naturalists" are not naturalists on Harman's reading, for being a naturalist in his sense means taking seriously the question of how moral properties are "located"; and "when one takes this attempt seriously, one will tend to become skeptical or relativistic" (p.34). Realism is not an option for naturalists, as far as Harman can see, for he takes realists to be committed to an absolutism that cannot account for empirical evidence for moral diversity.

Harman is correct in contrasting relativism and absolutism, and while many realists, including perhaps Cornell realists, are absolutists, there are most certainly forms of realism consistent with the denial of absolutism. There are ways other than

absolutism to capture the so-called "objectivity of morality." Harman seems to assume that conventions (cultures) are the only possible metaphysical grounds for the different systems of morality to be relativized to. If moral realism were true, the argument goes, then there must be one set of absolute moral facts, but if this were the case then all disagreements would be in principle reconcilable, there would be no way to account for the great disparity of actual moral practices, etc., etc.

Physical reality, however, is itself is not absolute but nonetheless deserves a realist's ontology (or at least we are assuming this in strictly metaethical debate).¹³ And, moreover, whenever we find disciplines that essentially demand attention to particular cases, we often find that there are no absolute answers to the pertinent questions. Consider medicine and nutrition, navigation, and engineering. In none of these disciplines are there absolute rules for what to do in all possible situations. To take medicine alone, human biology (and our individual biological make-ups) is different from, say, feline or canine biology and, thus, treating medical problems must be relativized to accommodate these differences. It would be foolish to conclude from this sort of relativization that medicine was not grounded upon facts that merit a realist's treatment. One may certainly embrace realism while denying absolutism. One need only think that moral situations obtain among unique individuals (or groups thereof), and that the facts which make each person unique must be considered in figuring out what ought to be done. If it is possible that changing the people involved *ipso facto* changes what ought to be done, then any absolutism demanding a single solution to all tokens of a type of a problem is facile at best. Morality is case sensitive in a way that is compatible with realism and incompatible with absolutism.

There is no reason to suppose that the shape of every human's happiness or *eudaimonia* will be the same; we are all individuals, each with our own talents and abilities, and what is in fact good for you may not in fact be good for me.¹⁴ Even if there are basic goods that we all must have some of, our particular needs will be tailored to facts about ourselves. It is important to learn that the facts which make us individuals are not necessarily morally irrelevant: to take Sartre's example, in trying to decide whether one should join the revolution or stay with one's sick mother, the facts concerning one's fitness to fight and nurse are certainly relevant. The best way, however, to understand how morality's realism and its relativity are consistent is to learn our lesson from advances in normative ethical theory of the last twenty years.

If there have been any advances in normative ethics in these last years, it is that traditional formulations of both consequentialism and deontology are unsatisfactory, for both assume a hard edged impartiality that, if lived by, alienates human beings from what makes for a good life.¹⁵ This has impact for us at both the personal and the species-wide level. At the personal level, we find that our own individual values are relativized so that, for example, as individuals we do not absolutely value all humans equally; in certain circumstances, it is acceptable to give preferential treatment to those we love over strangers. At the species wide level, we are concerned with a morality that respects the implications of the fact that we are *Homo Sapiens* with certain needs, talents, and natures. These natural facts are not morally inert. When studying epistemology, our theories are concerned with the nature of justification in general, and not with the specific ways that humans may go about justifying or being justified. Morality is different.

Here, we want our theories to be tailored to fit the facts of a specifically human form of life, for without this they will cease to be practical. There is no reason to suppose (*contra* Kant and other absolutists) that what makes for a good human life will be what makes for the good life of any rational creature. If there are other intelligent or rational forms of life in the universe, we cannot assume that there is some necessary overlap between their good and ours.

How does this bear on the semantic problem at hand? Slipping into the thought experiment again, we have encountered Twin Earthlings who have a different morality than ours; they predicate "good" differently than we do. The realist should respond like this: either there is a relevant difference between us and the Twin Earthlings that accounts for the difference in our moral theories and predications (and no one has made any mistakes, even though we don't all give the same answers), or there is no such difference and either we've gotten it wrong or they have. Or, of course, it is also possible that both us and the Twin Earthlings have gotten it wrong, just as we have been wrong about many empirical matters in the past.

This response is an adequate solution to the semantic problem. It is not, however, consistent with the ontologies of the most familiar forms of moral realism. The Cornell realists have likened the property of being good to the property of being a certain kind of stuff (as water is a certain kind of stuff).¹⁶ Michael Smith says that doing the right thing is doing what a perfectly rational, fully informed being would desire to do (cf. note 4). Secondary quality theorists say that the property of being good is rigidly regulated by the actual judging practices of certain observers in certain conditions. None of these are amenable to the kind of relativism just discussed as the way out of the problem, for each gets

tripped up by their *identifications* of goodness with some particular (set of) properties, such that what is identified with goodness at the actual world is thus necessarily identical with goodness at all possible worlds. This is a strong form of absolutism. None of these ontologies can admit the possibility that both us and the Twin Earthlings have not made any errors, even though we do not agree on all matters of morality.

What we need is a term or a sort of term which can serve as a semantic model for "goodness"; "water" has been tried, "milk" has been tried, "redness" has been tried and these all engender problems when it comes to Moral Twin Earth. The problem with all these terms is that they do not reflect the theoretical aspect of our concept of goodness or the property of *goodness*. "Rational" has been tried, but rationality has a sort of rigidity or absoluteness about it that also causes problems at Moral Twin Earth. What will work as a semantic model for "goodness" is "healthiness." "Healthiness" is appropriately theoretical and is understood to be concerned with proper function, but this does not make it a "functional" property; there is no "functional role" of "healthiness" that can be specified or captured by a definite description. Primarily, "healthiness" names a property of living objects, so that being healthy is best understood as a *way* of being alive or a *way* of functioning (just as being sick is a way of being alive). It is a property of relations which obtain among subsystems. While a little more will need to be said about the metaphysics of health, our first concern is with the semantics of "health" and the place to begin is by looking to what has been said in the past about it.¹⁷

While Aristotle discusses different ways of being healthy, Aquinas explicitly addressed the semantic structure of "healthiness" to help him understand some aspects of theological discourse. He saw that the

various uses of "health" can serve as an example of one kind of analogical predication, which he then used to help understand discourse about God's nature.¹⁸ We may leave aside theological discourse, however, for our concerns are with analogical predication and "health" *per se*. Aquinas thought something like this: the basic or primary predication of "healthy" is to organisms, organs, and tissues, but "healthy" admits of two other kinds of analogical predication as well. We also call "healthy" causes as well as signs of an organism's "healthiness" (in the primary sense). For the sake of convenience, let's call these three predicates "b-health" for the basic predicate, "c-health" for causes of health, and "s-health" for signs of health. We may then say that food is c-healthy if it promotes or causes one to be b-healthy, and elastic skin tone or blood pressure within a certain range is s-healthy if these turn out to be signs, symptoms, or indications of b-health.

How would such semantic complications arise? Well, to begin with, b-health, as a property, is instantiated by living or biological tissues, organs, or organisms when they are functioning properly. As mentioned above, in the broadest and most metaphysical terms, b-health is a *way of being*. B-health is the efficient living actualization of inherent potential energy within matter: a living item is "b-healthy" if it is "functioning properly," where the "properly" really is of the last significance (see note 17 above). A function may be either functioning well or poorly, properly or it may be malfunctioning. It is important to recognize that this does not make the property of b-health a functional property: it is best understood as a property of a function. (Malfunctioning and proper functioning will have the same status, and if the malfunctioning of a function is not itself a functional property, then neither is

proper functioning.) The property b-health is instantiated by living systems and is *a way of being alive*. When a function is a living function, as the function of a heart is a living function, and it is properly functioning, b-healthiness is instantiated.

At some early point, we must have posited the existence of the property of b-health, as, at the very least, the name for how it is in the absence of illness or death. Since we do not empirically observe [b-]healthiness as we do water, or redness, and must thereby “posit” it in some theoretical sense, perhaps there is some sense in saying that “[b-]health” is a “theoretical term.” But importantly, for the reasons just mentioned, it is not a term that can be learned through the causal theory of reference: there are no *healthons* out there to be discovered, the more of which we have the healthier we are. Nor is “[b-]health” a term that can be considered a rigid designator: there is no one set way organs or organisms must be in order for b-healthiness to be instantiated at a world. There is no way to give a full and general description of the truth conditions which would warrant a predication of “[b-]healthiness”; even though there is a sense in which it is a theoretical term, one cannot define it in terms of b-health’s functional role as some (like Armstrong and Lewis) that think one can describe the functional role of “pain.” “[B-]healthiness” cannot be defined by the Ramsey-Carnap-Lewis method of defining theoretical terms. The only way to perhaps “capture” its content in a description would be to serially describe what counts as “healthiness” as it is particularized to every individual at every time at every world. This would, however, be a vacuous or trivial “capture”: one does not succeed in defining a theoretical term by describing serially each set of truth conditions that would warrant a predication of that term. This is most especially so if one cannot say

anything more general regarding what it is about *these* sets of truth conditions which warrants the predication of “[b-]healthiness.” If this is not clear, try to imagine what it would be Ramsify a definition for “living” and being given a list describing all the living things in all possible worlds. The discussion devolves into something no more helpful than talk about the dormative effect of opium.

This is not, of course, the full story of how “[b-]health” gained meaning or reference, but such a story is not required for seeing how it may serve as the beginnings of a model for “good.” For now, we may say that the semantics of “[b-]health” can be understood in terms of being guided by the presence of the property b-health as instantiated by tissues, organs, or organisms. It is found in propositions such as “It [He, She] is healthy.” B-health may also be predicated over groups of individuals, where it may refer to a statistical norm of health of the group as a whole: “It’s a [b-]healthy herd of cattle.” Since it seems that straightforward definition is impossible, let it suffice to say that, while bordering on metaphor, tissues, organs, and organisms are “b-healthy” when they are “robust,” “strong,” “flourishing,” or “doing well.” A “b-healthy” body or organ is one which is “doing what it ought to do,” in the sense of its having a biological “job” or “function” to perform and “doing so proficiently,” “efficiently,” or “properly.” Perhaps it suffices for now to simply say that when an object is “b-healthy,” it is “functioning properly,” or “doing what it ought to do.”

The semantics of “c-health” are perhaps easy to grasp, due to the existence of antonyms that give fairly clear insight into its meaning. If “c-health” is predicated over causes of “b-health,” then its antonym would be either “toxicity” or “poisonousness.” It is found in propositions such as

"chicken soup is [c-]healthy food" or "mercury is poisonous." Now, here we have a property upon which a realist can readily find purchase; no one with the strength of their convictions can survive while denying the reality of the existence of toxic or poisonous substances. Thus, there are substances which are properly called "[c-]healthy." Unsurprisingly, a list of such substances is impossible to adduce, if only because of the relativity of "c-healthy" and "toxic." At the very least, these are relativized to species, such that for cows, eating grass is c-healthy while it is not for humans, and as well to certain individuals, such that iron rich diets are c-healthy for those with anemia are not for those who lack it. Interestingly, what counts as "c-healthy" or "toxic" can even depend on amounts and situation, such that large amounts of arsenic are toxic for humans while in certain instances of sickness (b-unhealthiness) small amounts are used as c-healthy medicine. While an antonym for substances which are "c-healthy" is "toxic," there are behaviors which can be called "c-healthy," such as exercise, but there do not seem to be behaviors which can be called "toxic." Being slothful is not toxic for a person, but it is not c-healthy either. Sloth can lead to the build up of toxic substances that would otherwise be purged through exercise, but not exercising is still not properly called "toxic." (Smoking cigarettes might be considered a "toxic" behavior, but this is in virtue of the toxic substances contained in tobacco.) The same may be said of other lacks of what is c-healthy: a prolonged lack of vitamin C causes scurvy, but not eating vitamin C is not properly called "toxic," even though consciously not doing so will kill a person. In the same way, not eating mercury is not properly called "c-healthy." It may be the case that some of these semantic intuitions about what to call absences

of "c-health" or "toxicity" are not generally shared by all English speakers, but this is not quite to the point. What is most important is to note that we have at least some clear notions about when to predicate "c-health" and when to predicate "toxic": whatever further debates exist, they rest on a shared understanding of some central cases, and disagreements about other cases are thus substantial and not merely semantic.

In our third use of "health," we find it being predicated of signs or symptoms of b-health, and in propositions such as "you have a healthy complexion." We get some insight in the kind of situation in which this term is used by attending to what Grice says about the difference between "natural" and "non-natural meaning."¹⁹ The former is at play when we say "those spots mean measles" or "the tree having x number of rings means that it is x years old." Neither of these propositions can be restated by following the verb "mean" with a phrase in inverted quotes. (E.g., "those spots mean 'measles'" makes no sense.) Non-natural meaning is at play when saying "those three rings on the bell (of the bus) means 'the bus is full'" or "he is telling you he loves you when he says 'be careful.'" As far as non-natural meaning goes, there is no commitment here to a Gricean theory of semantic or non-natural meaning. Our concern is with natural meaning. So, to begin with we may say that when x is a sign of y 's b-health, we may say that x naturally means that y is b-healthy. If x naturally means y is b-healthy, then x is s-healthy. If x is the elasticity of y 's b-healthy skin, a doctor might say to a medical student, "elasticity such as this is [s-]healthy." The relationship between y 's b-health and a sign thereof is most often a causal relationship, as in case of spots and measles, but it may also be a merely logical relationship as we note in "the leg's being bent like that means that it is broken"

(and there is only a logical distinction to be had between the bend and the break). As a predicate, “[s-]healthy” gains its (natural) meaning by referring to signs of b-health (e.g., a healthy complexion). And as a predicate, “s-health” (like “b-” and “c-health”) has an indefinitely large extension; this is the case given that every individual’s b-health will itself be individualized, and thus the signs of every individual’s b-health will also be individualized. Two individual conspecifics might have the same blood pressure, where for one this would (naturally) mean that the cardio-vascular system is b-healthy, while for another this might mean there is a great risk of imminent heart attack.

Thus, the basic structure of the semantics of “healthiness” has a rather complex shape that Aquinas has at least roughed out. The complications themselves are both unfortunate and fortunate: unfortunate insofar as they make understanding the semantics of “healthiness” difficult, fortunate because of the fact that the semantics of “goodness” have the same complex shape. And because this shape is as idiosyncratic as it is, it is proportionally hard to see as mere coincidence the similarities between the semantics of “healthiness” and “goodness.” (As a *prima facie* start, note how each handles negation, has antonyms, etc.)

But perhaps it is already obvious that there are correlates of “b-,” “c-,” and “s-health” in the semantics of “goodness” (and “badness”). These have already been discussed above in terms of the three ways that “good” is predicated. (The deck was a bit stacked.) We first predicate “good” of people or lives thereof, and this is the correlate of “b-health.” And the distinction among uses of “good” discernible by their taking different prepositions is isomorphic to the distinction between “c-health” and “s-health.” We commonly say both “tofu is healthy *for* a person” as we say that “doing charity is good *for* a person”; we also

say “it is a sign *of* good health to have low blood pressure” as we say “it is a sign *of* a good person to do care for one’s aging parents.” It will be worthwhile to spend a few of paragraphs laying out a few of the details (signs) of the structural isomorphism between the semantics of “health” and “good.”

Beginning with b-health and virtue ethics, and in terms of the most basic and agreed upon of our moral data, given the perspective of time and the coolness of reflective judgment, we think that Socrates, Jesus, and the Buddha were in all probability good people, while Thrasymachus, Hitler, and Stalin were bad. One needn’t agree with the list; the important point is that the predicates “good” and “bad,” used in this way, refer to properties of people: the properties of being good or bad are instantiated by people, or perhaps more precisely either by people’s lives or people’s characters. When we talk about “good people” or someone living a “good life,” we are using a sense of “good” that is the correlate of “b-health.” And when philosophers focus on this agent-centered use of “good,” they end up talking about virtue theory or eudaimonistic ethics: here the goal of living morally is to live a “b-good” life, or to develop a “b-good” character. We must remember, however, the relativization of goodness that was discussed above: what counts as a b-good life will be relativized to different species, members of which possess the same general traits, and even within a species, b-goodness will be relativized to individuals, tailored to each person’s particularized talents and abilities. And these various sorts of relativization are mirrored by our understanding of b-health as “proper functioning” which itself is relativized to the constitution of any particular individual’s b-health and will thus be similarly tailored or relativized to that individual.

But we also predicate “goodness” upon two kinds of conduct, act or performance.

To begin, we predicate “goodness” of conduct that we take to promote ends like happiness or *eudaimonia* or pleasure. And when we say that “good acts” are those that cause or promote these ends, we are justifying the goodness of that conduct by pointing to its consequences. And despite the considerable problems that consequentialist thought has, it seems undeniable that in common moral language there are certain times when we do value acts sheerly for their results and the manner or motive by which these results are brought about is irrelevant. Here, is the correlate of “c-health.” When moral discourse centers on calling certain acts “good” in virtue of what they produce or what they cause, these are c-good. Just as certain foods that causally engender certain (functionally or biologically) desirable states of affairs are called “healthy” (“c-healthy”), so too may performing certain acts instrumentally engender certain morally desirable ends. If one were to realize one’s own selfishness, one might do some charity work, as it might be good *for* one. We say to truculent children, “It doesn’t matter whether you want to visit your relatives or not, it is good *for* you, and you will go.” Focusing on those acts which we take to be beneficial to the production of certain moral ends, we find ourselves engaged in the consequentialist’s dialectic: acts are called “c-good *for*” what they accomplish, or as means to ends.

But there is another way to predicate “goodness” of conduct. For Kant begins his *Groundwork* by saying that the only thing that is [b-]good in itself, in some unspoilable way, is a good will. However, for reasons relating to the philosophy of mind and the rejection of “faculties,” it might be wrong to hypostatize the will: perhaps the right way to understand a good will may be to think of it as being manifest only in actions of a certain sort or with a certain kind of motive. So, more or less according

to Kant, these are actions that we do not value because of what they promote, but they gain their value from their source or motivation. These acts are “good” if they are motivated by respect for duty, or love of God or humanity or the kingdom of ends, or however the sources of normativity are formulated. This is the deontologist’s dialectic, and when we are engaged in it we are using the word “good” in a way analogous to “s-health.” To see why this is so it will be helpful to reflect on an example of an act with deontological value.²⁰ Suppose that the Buddha and Socrates are best friends and that the Buddha winds up in a hospital. If Socrates visits the Buddha, the visit may be motivated by certain ends: it might have c-good consequences for the Buddha’s health, or Socrates might make the visit in order to avoid feeling guilty for having not gone. On the other hand, we ordinarily say that it would be best for Socrates to go, not because of anything to be accomplished by going, but because he is the Buddha’s best friend and best friends visit each other in the hospital. It would be “[s-]good *of*” Socrates to go. Put perhaps perspicuously, or at least in a more clearly Kantian cadence: members of the kingdom of ends visit each other in the hospital. The act has value, not for what it can do or accomplish, but for what it is; and what it is, in Kantian terms, is an expression *of* a friend’s b-good will, and (for Kant at least) the only way for conduct to have moral value is if it is an expression of a b-good will. It is an act that is done *from* the proper motivations. We have learned that such acts are how a b-good will performs or expresses itself, acts of this sort are evidence of a b-good will. Such acts are simply how a member of the kingdom of ends behaves, and when we observe someone behaving like a member of the kingdom of ends, we take this as a evidence of, as a sign of that person’s b-goodness. Part of how we judge

a person morally (or ourselves morally), part of how we make a diagnosis of a person's moral character, is by learning how to detect signs or expressions of a b-good will. Doing volunteer charity work is a defeasible sign of a b-good person (though someone might be doing it for ulterior motives), and this is just the way that we use signs like being able to run a marathon in judging a person's health. Thus, acts are s-good, insofar as they are non-instrumental signs, indications or expressions of a b-good will. If we wish to de-hypostatize "the will," we may instead talk in terms of motivation: acts are s-good insofar as they have the correct non-instrumental motivations, where these are something like acting *from* a respect for friends, duty or God, or love of humanity, etc.. Insofar as we predicate "[s-]good" of certain conduct for non-instrumental reasons, we do so for what this conduct is an expression of. If an act's etiology has its genesis in a person's b-good will, then it is an s-good act. Just as we take it as a mortally important sign of (ill-)health that one is out of breath after climbing a single set of stairs, because of what this tells us about the b-health of one's cardio-vascular system, we attach importance to certain conduct by calling it "good," because of what it tells us about an agent's moral character or b-good will.

In summary, we have detected a variety of signs or indications that the semantics of how we actually talk about morality, and in particular the discourse of normative ethics, can be modeled on the tri-partite semantics of health: just as "health" has three forms of predication, so too does "good." Upon a moment's reflection, it should be no surprise that the debate between consequentialists and deontologists has been irresolvable: an examination of ordinary moral discourse shows that acts are called "good" for more than one reason,

and this is reflected in the debates of moral normative theorizing.²¹ Obviously, ordinary moral discourse allows us to speak meaningfully as both consequentialists and deontologists; if this were not the case, the debate would never have gotten off the ground. We speak in different ways when we are valuing conduct for different reasons, and both ways of speaking are quite legitimate. At the very least we have found a large coincidence: we could have guessed that the semantics of "goodness" and "healthiness" are both very complicated, and this makes it all the more intriguing to find that their complications are similar in so many ways. The degree to which it is improbable that this large coincidence is merely coincidence, is the degree of credence which accrues to the thesis that goodness and healthiness share their ontic status.

To lend this thesis even further credence, we can finish by noting that this form of moral realism is immune to Horgan and Timmons's "new wave" semantic problem. Perhaps it should first be noted that the Twin Earth experiment could not go as it is originally described. The original thought experiment was based on the idea that we could discover which properties moral discourse is attempting to track and thereby find ourselves speaking as either deontologists or consequentialists (as the case may be). We now find that everyday moral discourse employs constructions, like "good character," "good for," and "good of" which indicate that our actual discourse tracks many sorts of properties. Without eradicating large chunks of ordinary moral discourse, we could not find ourselves as deontologists who differ from consequentialist Twin Earthlings only in the particular acts to which we predicate the word "good." We will not merely disagree with them about cases, for, if we have learned anything, it is that deontologists and consequentialists speak in different

ways; both theories employ their own distinctive semantics for how "good" is used. Earthlings and Twin Earthlings, if they were limited to the resources of one theory or the other, would have moral discourses with very different semantic structures. And this is quite different than in the case of "water," where the semantics used to discuss H_2O and XYZ will be identical. This may cause one to question the justifiability of using the Twin Earth scenario which demands that as much English as possible be the same in both places.

Be that as it may, let us make as much of a problem as we can for the model of "good" based on "health" and see how it stands up to Moral Twin Earth. Let's do so by first seeing what we should say about "health" at a world we can call *Medical Twin Earth*. We find ourselves on Medical Twin Earth and note that adult Twin Earthlings there wear little heavy clothing in the winter, and have no hot water heaters anywhere. It turns out that on Medical Twin Earth they have "discovered" that warmth, while certainly pleasant, weakens the adult cardio-vascular and immune systems. They think warmth is c-unhealthy for b-healthy adults, while we disagree. The first thing we must figure out is if Twin Earthling physiology is indeed the same as ours, for if there are relevant differences between us, then none of us may be making mistakes when saying that something that is c-healthy for us is c-unhealthy for them. This is only the sort of relativism which our discussion above showed to be compatible with realism. If there are no differences, then Earthlings and Twin Earthlings cannot both be right. Figuring out who has made the error will be complicated and difficult work, forcing us to catalogue all our signs of b-health, investigating empirically the life-long effects of increased coldness on our b-health, doing the complicated epidemiological statistics

that all this data will produce, etc., etc.. In fact, we might not be able to find out who is making the mistake for there might be disagreements over how to conduct the empirical investigation. (Consider the extent disagreements about how much cholesterol is healthy.) On the other hand, we might not find any substantial difference in health between those who keep cold and those that do not; there may be no functional difference between the two. But we certainly do not conclude from any of this that we ought to be realists about health because of these disagreements.

Now take the case with morality. Let's focus on a case similar to one discussed by Horgan and Timmons.²² Here on Earth, based on our acceptance of a community minded utilitarianism, we decide that some activist welfare safety net is a sign of a b-good moral community, while on Moral Twin Earth they have decided, based on an individually minded deontology, that living in a libertarian community is best. Thus, we speak differently from them, disagreeing with them (only?) in particular cases of how the word "good" ought to be used. We proceed as we just did on Medical Twin Earth. First, we must look to see if there are relevant differences between us and them; perhaps these Twin Earthlings have psycho-biologically evolved in such a way that they are creatures best off alone (imagine finding out that a substantial percentage of Twin Earthlings live as hermits), whereas we on Earth are more clearly "social animals." If this were the case, then maybe no one has made any mistakes at all about what is morally good. If, however, there are no relevant differences, then we must examine the states of affairs in both worlds to see which system has left its inhabitants with happier, more fulfilling lives. Obviously, this empirical research will be complicated and perhaps there will be disagreements over how it

should be done. Maybe we could figure out which moral system is better, maybe not. Maybe we find that the best system is some mixture of the two; that adopting one system to the complete exclusion of the other is the source of the problem. Maybe we find that everyone's practices are riddled with errors and that we really need to start over. We might not be able to determine which moral system is better. What is crucial is that none of these results are problematic for the theory at hand. Inter-planetary moral disagreements can genuinely be engaged without each side necessarily begging the question against the other by assuming that their theory is the true one. We do not talk past one another, and this is just what we should expect from a successful semantics for moral terms.

It is important and yet perhaps difficult to see how the Moral Twin Earth problem fails to arise for "good" as presently understood. As a helpful analogy, consider the following variation on the Medical Twin Earth case. First, assume that Earthlings and Twin Earthlings are functionally equivalent from a biological point of view. Given this, imagine that everyone on Earth has completely converged on Western style allopathic medical practices, and when we get to Twin Earth we find them talking about something they call "chi" and basing all their medicine on sticking small pins into people and giving them herbs. Imagine that on Twin Earth they have com-

pletely converged upon what we (in the West) actually call Eastern style "Chinese medicine." Earthlings and Twin Earthlings will then have roughly similar conceptions of what counts as b-health, though we may explain the reasons for this differently. We will agree often about what is s-healthy and c-healthy, in the easy cases, but there will be disagreement over a very wide variety of (hard) cases. In this scenario, it might be impossible to achieve a convergence between our theory and theirs. (Perhaps much like the way it may be impossible to get a convergence between creationists and evolutionary theorists.) It will at least be very difficult for us to determine which style of medicine is best for creatures like us, and is it not likely that we should end up adopting fully one style of medicine to the exclusion of the other. None of the difficulties we have in discerning the truth about medicine make us think that we should be irrealists about health; Earthlings and Twin Earthlings need not "talk past each other" regarding health and indeed we all could certainly have genuine disagreements. It all falls out exactly as one would expect. If we model "goodness" upon "healthiness" and Medical Twin Earth is not a problem for the semantics of "healthiness," then Moral Twin Earth does not present any problems for the rules of "goodness" as laid out above.²³

University of Connecticut, Storrs

NOTES

1. The "Cornell realists" do not all share a single theory, but all have similar ontologies for goodness and semantics for "goodness." See, for example, Nicholas Sturgeon's "Moral Explanations," Richard Boyd's "How to Be a Moral Realist," and Geoffrey Sayre-McCord's "Moral Theory and Explanatory Impotence," all reprinted in *Essays on Moral Realism*, ed. G. Sayre-McCord (Ithaca, N.Y.: Cornell University Press, 1988); David Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989); and Peter Railton, "Moral Realism," *The Philosophical Review*, vol. 95 (1986), pp. 163–207.
2. Hilary Putnam, *Reason, Truth, and History* (Cambridge: Cambridge University Press, 1981); Saul Kripke, "Naming and Necessity," *Semantics of Natural Language*, ed. D. Davidson and G. Harman (Dordrecht: Reidel, 1972).
3. Terry Horgan and Mark Timmons, "New Wave Moral Realism Meets Moral Twin Earth," in *Rationality, Morality, and Self-Interest*, ed. J. Heil (Lanham, Md.: Rowman & Littlefield Publishers, 1993), pp. 115–133. (In my opinion this is the most general treatment of the problem.) "New Wave Moral Realism Meets Moral Twin Earth," *Journal for Philosophical Research*, vol. 16 (1991), pp. 447–465; "Troubles on Moral Twin Earth: Moral Queerness Revived," *Synthese*, vol. 92 (1992), pp. 221–260; "From Moral Realism to Moral Relativism in One Easy Step," *Critica*, vol. 28, no. 83 (1996), pp. 3–39; "Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived," *Philosophical Papers*, vol. 21, no. 3 (1992), pp. 153–175; "Troubles for Michael Smith's Metaethical Rationalism," *Philosophical Papers*, vol. 25, no. 3 (1996), pp. 203–231; and "Copping Out on Moral Twin Earth," *Synthese*, vol. 124 (2000), pp. 139–152.
4. Versions of these theses are found throughout the series of papers listed above in footnote 3. These are taken from Horgan and Timmons, "New Wave Moral Realism Meets Moral Twin Earth," in *Rationality, Morality, and Self-Interest*, p. 123. One might wonder what work the phrase "by human beings" is doing here, especially when there is no single term *t* used by human beings to refer to instantiations of goodness. Presumably, we are to think of *t* as a set of synonymous terms each used in its own language to refer to instantiations of goodness. Note how this builds absolutism into the picture. *CRT* and *CSN* can be reformulated to avoid this problem: replace "by human beings" with "by speakers of the language of which *t* is a part."
5. For a conclusion to the thought experiment different than Putnam's, see George Bealer's "Limits of Scientific Essentialism," *Philosophical Perspectives*, vol. 1 (1988), pp. 289–364.
Bealer here suggests that "water" is a functional term, so that if XYZ plays the same role in the lives of Twin Earthlings as H₂O plays in our own, then we should be willing to say that "water" has the same meaning on both planets. It will become clear, below, that, regardless of the success of Bealer's gambit in the case of Putnam's essentialism, it is a powerful move to employ in the moral context.
6. For more on this contrast, see Robert B. Loudon, *Morality and Moral Theory* (New York: Oxford University Press), 1992.
7. John Rawls, "Symposium: Justice as Fairness," *The Journal of Philosophy*, vol. 54, no. 22 (Oct. 1957), pp. 653–662; Barbara Herman, "On the Value of Acting from the Motive of Duty," *Philosophical Review*, vol. 90, no. 3 (July 1981).
8. Horgan and Timmons, "New Wave Moral Realism Meets Moral Twin Earth," in *Rationality, Morality, and Self-Interest*, p. 133.
9. Michael Smith, *The Moral Problem* (Oxford: Basil Blackwell), 1994. David Copp, "Milk, Honey, and the Good Life on Moral Twin Earth," *Synthese*, vol. 124 (2000), pp. 113–137; Frank Jackson,

From Metaphysics to Ethics (Oxford: Clarendon Press, 1998). For the application of the recipe to Smith and Copp, see the citations above at note 3. In private communication, Horgan and Timmons report that they can “cook-up” a version of their recipe to show where Jackson has gone wrong.

10. For the application of the Moral Twin problem to secondary quality analyses of value, see Sean Holland, “Dispositional Theories of Value Meet Moral Twin Earth,” *American Philosophical Quarterly*, vol. 38, no. 2 (April 2001) p. 177–195. For the secondary quality analyses themselves, see, for example, John McDowell, “Aesthetic Value, Objectivity, and the Fabric of the World,” in *Pleasure, Preference, and Value*, ed. E. Schaper (Cambridge: Cambridge University Press, 1983), pp. 1–16; and McDowell’s “Values and Secondary Qualities,” in *Morality and Objectivity*, ed. Ted Honderich (London: Routledge and Kegan Paul, 1988), pp. 110–129; David Wiggins, “A Sensible Subjectivism,” in *Needs, Values, Truth* (Oxford: Basil Blackwell), 1991. David Lewis, Mark Johnston, and Michael Smith all have papers entitled “Dispositional Theories of Value,” all printed in *Proceedings of the Aristotelian Society*, supp. vol. 63 (1989).

11. J. L. Mackie, *Ethics* (Harmondsworth: Penguin Books, 1977), pp. 36–38.

12. Gilbert Harman, “Is There a Single True Morality?” in *Morality, Reason, and Truth*, ed. D. Copp and D. Zimmerman (Rowman and Allanheld, 1985), pp. 27–48.

13. Consider, for example, Einstein’s account of time. He is obviously no absolutist, but his empirical account about the nature of time cannot rightly be called any sort of “non-realism.” He clearly believes that time is discovered and is neither conventional nor invented, that it operates according to natural law, and while it may be relativized to frames of reference, this in no way makes it somehow unfit for realism. Entropy is also a relativized aspect of physical reality (see Lawrence Sklar, *Physics and Chance* [Cambridge: Cambridge University Press, 1993], p. 346).

14. Even Kant, our paradigm of absolutism, recognizes that “a metaphysics of morals cannot be based upon anthropology but can still be applied to it.” *The Metaphysics of Morals*, trans. Mary Gregor (Cambridge: Cambridge University Press, [1797] 1996), p. 10.

This may sound like it is at odds with the so-called “universalizability” of moral language, but it is not. The question concerns the scope of the universalizability. If we fix the scope to the level of functional equivalence, then we may preserve universalizability by saying that what one person ought to do in a given case is what anyone ought to do in a functionally equivalent case. This is perfectly compatible with a relativization of morality, namely to all functionally equivalent cases, and to realism as well. For more on universalizability and its proper scope, see my “Prescriptions are Assertions: An Essay on Moral Syntax,” *American Philosophical Quarterly*, vol. 35, no. 1 (1998), pp. 1–20.

15. See, for example, Michael Stocker, “The Schizophrenia of Modern Ethical Theories,” *The Journal of Philosophy*, vol. 73 (1976), pp. 453–466; Laurence Thomas, *Living Morally* (Philadelphia: Temple University Press, 1989); Bernard Williams, *Moral Luck* (Oxford: Oxford University Press, 1981); Susan Wolf, “Moral Saints,” *Journal of Philosophy*, vol. 79 (1982), pp. 419–439.

Of course, modern theorists have worked to make consequentialism and deontology more humane. See, for example, Peter Railton’s “Alienation, Consequentialism, and the Demands of Morality,” *Philosophy and Public Affairs*, vol. 13, no. 2 (Spring 1984); and again, see Barbara Herman’s “On the Value of Acting from the Motive of Duty,” *Philosophical Review*, vol. 90, no. 3 (July 1981), pp. 359–382.

16. One might think that moral properties might be ontologically of the same sort as mental properties. And this is indeed at the heart of Brink’s version of moral realism (*Moral Realism and the Foundations of Ethics* [Cambridge: Cambridge University Press, 1989]). But one cannot elucidate an ontology for moral properties by comparison to mental properties, for the nature and

ontological status of the mind is itself incredibly obscure and poorly understood. If eliminative materialism turns out to be the truth about the mental realm, would that entail a similar result for the moral realm? This would be unfortunate for those who want to be moral realists, who will fare better by staying away from the conundrums of the mind when looking for insight.

17. In what follows, I will be talking a lot about proper function and proper functioning. There are many different possible senses for this term. I use the adjective “proper” and the adverb “properly” as ways of modifying the noun “function” and the verb “to function,” respectively. I think this is the original and common sense way of discussing functions, biological and artificial. It is the sense of “function” at play when we say that a biological heart and an artificial heart may be “functionally equivalent”; it is also the sense of function what relates “proper function” to “malfunction” as antonyms. So, as a contrast, this is not the sense of “proper function” which has been discussed by Ruth Millikan, for example, in *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press, 1984).

There is a small literature on health and proper biological function. See for example, Christopher Boorse, “Health as a Theoretical Concept,” *Philosophy of Science*, vol. 44 (1977), pp. 542–573; Ruth Macklin, “Mental Health and Mental Illness,” *Philosophy of Science* vol. 39, no. 3 (1972), pp. 341–365. For treatment of the relations between healthiness and goodness, see my *Moral Reality* (New York: Oxford University Press, 2001).

18. The basic idea is that, when God says, “I know,” what he means by “knowing” is not what humans mean when they use the word. Nevertheless, the meaning of God’s “knowledge” is analogically related to human “knowledge.” I here follow mainly William Alston, “Aquinas on Theological Predication,” in *Reasoned Faith*, ed. E. Stump (Ithaca, N.Y.: Cornell University Press, 1989).

19. H. P. Grice, “Meaning,” *Philosophical Review*, vol. 66 (1957), pp. 377–388.

20. This example is adapted from one used by Michael Stocker in his paper “The Schizophrenia of Modern Moral Theories,” *The Journal of Philosophy*, vol. 73, pp. 453–466.

21. Interestingly, if “good acts” be so for either deontological or consequentialist reasons, then there might be a sense in which “good” is like “jade,” insofar as the latter refers to two different substances. This is not, however, a reason to revert to a semantic model for “good” based on substance terms or to speculate upon whether or not “good” is a natural kind term. The important point is that a moral semantics modeled on “health” gives us as much realism as the realist needs.

22. See Horgan and Timmons, “From Moral Realism to Moral Relativism in One Easy Step,” *op. cit.*

23. I would like to thank the members of the Department of Philosophy at the University of Memphis, to whom I presented an early ancestor of this paper. In particular, I would like to thank Terry Horgan and Mark Timmons for all their guidance. As usual, Michael P. Lynch gave me sage advice throughout the whole process.