

Neutral and Relative Value after Moore*

Michael Smith

We have inherited from G. E. Moore the following, I think attractive, analysis of what we ought to do in terms of what is good and bad.

$\forall x$ (x ought to Φ in certain circumstances C if and only if Φ -ing is that action, of the actions that x can perform in C , that produces the most good and the least bad).¹

The analysis proposes a reduction of one moral concept (the concept of what we ought to do) to another pair of moral concepts (the concepts of goodness and badness).

In order to fully understand the concept of what we ought to do, we must therefore combine the proposed analysis with an account of the nature of goodness and badness. Moore combines the analysis with his distinctive view of goodness and badness as simple, nonnatural properties. However, as we shall see, if the analysis is combined with this Moorean account of goodness and badness, then various ethical doctrines that mightn't have looked to be incoherent on the face of it turn

* An earlier version of this article was presented at "G. E. Moore's *Principia Ethica*: A Century Later," a conference organized by Kit Wellman and sponsored by the Jean Beer Blumenfeld Center for Ethics at Georgia State University. I am especially grateful for the many helpful suggestions I received from my commentator on that occasion, John Deigh, and also for fruitful discussions with the other participants. Thanks are also due to audiences at the University of Arizona at Tucson, Australian National University, University of Copenhagen, Lund University, University of Melbourne, Monash University, University of North Carolina at Chapel Hill, Oxford University, University of Singapore, and University of Sydney for their comments on a relative of the present article. John Broome, David Chalmers, Tyler Doggett, James Dreier, Andy Egan, Lloyd Humberstone, Thomas Hurka, Susan Hurley, Jenann Ismael, Frank Jackson, Jeanette Kennett, Philip Pettit, Wlodek Rabinowicz, Donald Regan, Andrew Reisner, Geoffrey Sayre-McCord, Samuel Scheffler, David Schmidtz, Daniel Stoljar, Jay Wallace, and Bernard Williams all made especially helpful remarks.

1. G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903). Moore expresses his subsequent misgivings in his "Reply to My Critics," in *The Philosophy of G. E. Moore*, ed. Paul Arthur Schilpp (Evanston, Ill.: Northwestern University Press, 1942), pp. 535–677.

Ethics 113 (April 2003): 576–598

© 2003 by The University of Chicago. All rights reserved. 0014-1704/2003/11303-0005\$10.00

out to be incoherent. For example, on the basis of this combination Moore himself goes on to argue that ethical egoism—the view that we are each obliged to maximize our own happiness—is self-contradictory.²

The interest of Moore's critique lies in the fact that, if successful, it would extend naturally to show the implausibility of all forms of what C. D. Broad calls "self-referential altruism": the view that we each have special obligations to our family members, our friends, our lovers, and so on.³ More generally, it would extend naturally to show the implausibility of all ethical theories which hold that there are relative values, as we would now put it—these are the values at stake in self-referentially altruistic conduct, the value of the weal and woe of our own family members, of our own friends, of our own lovers, and so on—in addition to, or perhaps rather than, neutral values: that is, the value of the weal and woe of people with whom we have no special relationship.⁴ When combined with the proposed analysis of obligation the upshot would thus be some form of consequentialism.⁵

Moore's conclusions thus plainly conflict with commonsense morality, as commonsense morality is replete with self-referentially altruistic elements. Indeed, commonsense morality seems to tell us that our obligations are a function of the relative weights of both the neutral and the relative values at stake in a particular choice situation. Suppose I am faced with a choice between giving a smaller benefit to my lover (a relative value) or a greater benefit to a complete stranger (a neutral value): imagine that the benefit is whatever can be bought for \$20, and, as it happens, my lover will benefit less from what she can buy for \$20

2. Here and throughout I assume, just for the sake of argument, that we should combine various doctrines about what we ought to do, including ethical egoism, with a hedonistic theory of value.

3. C. D. Broad, "Certain Features of Moore's Ethical Doctrines," in Schilpp, ed., pp. 41–67.

4. For the usage of the terms 'agent-relative' and 'agent-neutral' value I have in mind, see Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984). Part of the task of the present essay is to give a precise characterization of the distinction between agent-relative and agent-neutral values.

5. Like Parfit, I use 'consequentialism' to pick out that class of theories which holds not just that our obligations are a function of the goods and bads that result from our actions but, more narrowly, that they are a function of the neutral, as opposed to relative, goods and bads that result. Though this fits in with what I take to be current philosophical usage, it is worth emphasizing that there is therefore a broader class of theories that may be called "consequentialist" and that commonsense morality, at least as I go on to describe it in the text, is itself consequentialist in this broader sense. The fact that commonsense morality is consequentialist in this broader sense means that it avoids the main line of objection that Philip Pettit mounts against nonconsequentialism in his "Consequentialism," in *A Companion to Ethics*, ed. Peter Singer (Oxford: Blackwell, 1991), pp. 230–40. For useful discussion of the different senses in which theories may be said to be consequentialist, see James Dreier, "Structures of Normative Theories," *Monist* 76 (1993): 22–40.

than the stranger would benefit from what he could buy. Assuming that the consequences of our actions are otherwise identical, commonsense morality tells us that what we are obliged to do is a function of the differential in those two values. If the extra neutral value to be gained by giving the greater benefit to the stranger is less than the relative value that would be lost by failing to give the smaller benefit to my lover, then commonsense morality tells us that I have an obligation to give the smaller benefit to my lover, that is, to realize the relative value at stake. On the other hand, if the extra neutral value to be gained is greater than the relative value that would be lost, then commonsense morality tells us that I have an obligation to give the stranger that greater benefit. This is all radically at odds with what consequentialism tells us, which is that our obligations are a function of neutral values alone.

Of course, this is not to say that consequentialism isn't true. Perhaps commonsense morality is fundamentally mistaken, and, in the choice situation just described, we have an obligation to give the greater benefit to the stranger, much as consequentialism says, because there are only neutral values. For the record, I don't myself believe that this is so.⁶ But the issue on which I wish to focus here is not this but, rather, Moore's argument for the more ambitious claim that commonsense morality is not just false, but incoherent. That argument is potentially devastating. Yet, while some, like Broad, have tried to find fault with the argument, their attempts have not been entirely successful. Accordingly, my aim in the present article is to show where Moore's argument against ethical egoism, and thus his implicit argument against commonsense morality as well, goes wrong.

I begin by presenting Moore's own version of the argument (Sec. I) and Broad's response (Sec. II). Though, as we shall see, Broad's response is successful in its own terms, it leaves Moore's conclusion largely intact. After showing that Moore's argument readily generalizes to refute commonsense morality, if it succeeds at all (Sec. III), I suggest an alternative response to the argument (Sec. IV). I close by considering, and sketching replies to, three objections to that alternative response (Secs. V–VII).

I. MOORE'S ARGUMENT AGAINST ETHICAL EGOISM

Moore takes ethical egoism to be the view that

$\forall x (x \text{ ought to } \Phi \text{ in certain circumstances } C \text{ iff } \Phi \text{ing is that action, of the actions that } x \text{ can perform in } C, \text{ that produces the most good for } x \text{ and least bad for } x).$

6. Michael Smith, "Immodest Consequentialism and Character" ("Special Issue on Consequentialism and Character," ed. Julia Driver), *Utilitas* 13 (2001): 173–94.

The difference between ethical egoism and Moore's own view about the relationship between what we ought to do and what is good and bad is thus that ethical egoism substitutes talk of what is good for x for talk of what is good in an unqualified sense and, likewise, substitutes talk of what is bad for x for talk of what is bad in an unqualified sense.

Thus, according to ethical egoism, at least as Moore understands it, I ought to Φ in certain circumstances C if and only if Φ ing is the action, of the actions I can perform in C , that produces the most good for me and the least bad for me; you ought to Φ in circumstances C if and only if Φ ing is the action that you can perform that produces the most good for you and the least bad for you; and so on. By contrast, according to Moore's own view, I ought to Φ in C if and only if Φ ing is the action, of the actions I can perform in C , that produces the most good and least bad in an unqualified sense, and the same goes for you.

Moore's argument against ethical egoism takes issue with the substitution of "good for x " for "good." He asks what it could possibly mean for something to be good or bad for someone (from here on I will omit reference to what is bad just to simplify the discussion). His claim is that once we give the idea of something's being good for someone a plausible interpretation, it becomes clear that ethical egoism is self-contradictory.

What, then, is meant by "my own good"? In what sense can a thing be good *for me*? It is obvious, if we reflect, that the only thing which can belong to me, which can be mine, is something which is good, and not the fact that it is good. When, therefore, I talk of anything I get as "my own good," I must mean either that the thing I get is good, or that my possessing it is good. In both cases it is only the thing or the possession of it which is *mine*, and not *the goodness* of that thing or that possession. There is no longer any meaning in attaching the "my" to our predicate and saying: The possession of this *by me* is *my* good. . . . In short, when I talk of a thing as "my own good" all that I can mean is that something which will be exclusively mine, as my own pleasure is mine (whatever be the various senses of this relation denoted by "possession"), is *also good absolutely*; or rather that my possession of it is *good absolutely*. The good of it can in no possible sense be "private" or belong to me; any more than a thing can exist privately or for one person only. The only reason I can have for aiming at "my own good," is that it is *good absolutely* that what I so call should belong to me—*good absolutely* that I should have something, which, if I have it, others cannot have. But if it is good absolutely that I should have it, then everyone else has as much reason for aiming at *my* having it, as I have myself. If, therefore it is true of *any* single man's "interest" or "being happy" that it ought to be his sole ultimate end, this can only mean that that man's "interest" or "being happy" *is the*

sole good, the Universal Good, and the only thing that anybody ought to aim at. What Egoism holds, therefore, is that *each* man's being happy is the sole good—that a number of different things are *each* of them the only good thing there is—an absolute contradiction! No more complete and thorough refutation of any theory could be desired.⁷

As I see it, the argument Moore gives in this passage has two premises.

The first premise is that if being *F* is good for some person, *A*, then this entails that *A*'s possession of the property of being *F* is good in an unqualified sense. The second premise is that if *A*'s possession of the property of being *F* is good in an unqualified sense, then everyone, both *A* and everyone else, ought to do what they can to make *A* be *F*. Less schematically, given the assumption of hedonism, the first premise entails that when the ethical egoist says that being happy is good for *A*, what he thereby implies is that *A*'s being happy is good in an unqualified sense. But then, given the second premise, it follows that he is committed to claiming that everyone ought to do what they can to make *A* happy. This is why the ethical egoist commits himself to a contradiction when he goes on to say, as ethical egoists do, that what other people are obliged to do is not to make *A* happy but, rather, to make themselves happy. Let me consider these two premises in turn.

Moore's argument for the first premise is that since talk of something's being good for someone is talk of something that belongs to that person, and since it is incoherent to suppose that goodness itself, as distinct from the things that are good, could belong to someone, all that such talk can really amount to is the claim that the thing that is good belongs to that particular person. The crucial question is thus whether it really is incoherent to suppose that goodness itself, as distinct from the things that are good, could belong to particular people.

As I understand it, Moore thinks, correctly, that the property of goodness itself could belong to particular people only if goodness were a relational property, a property that relates the things that are good to those particular people. But, of course, as Moore sees things, goodness is not a relational property. Goodness is metaphysically simple. It is a simple, nonnatural property that belongs to the things that are good. The only thing that could possibly belong to people is thus, as Moore puts it, "something which is good"—*my* being happy, say, or *your* being happy—"and not the fact that it is good." The things that are good must therefore be understood to be, as he again puts it, "good *absolutely*," not as good in the sense of bearing some relation to a particular person.

The second premise in Moore's argument is the claim that "if it is good absolutely that I should have . . . [something] . . . , then

7. Moore, *Principia*, p. 99.

everyone else has as much reason for aiming at *my* having . . . [that thing] . . . , as I have myself.” Now, as I understand it, Moore’s reason for saying this is that it would be ad hoc to put a restriction on which of the various things that are good give rise to obligations. Suppose that my being happy and your being happy are both equally good—that is, that they both have the simple, nonnatural property of being good—and that our obligations are a function of the actions we can perform that promote things that are good. In that case the fact that your being happy is good could hardly put you under an obligation to make yourself happy, assuming that you are able to make yourself happy, without also putting me under an obligation to make you happy, assuming that I am able to make you happy. For there is no relevant difference between my being happy and your being happy in respect of goodness.

The upshot, as Moore sees things, is thus that ethical egoism is inconsistent with two plausible claims. Since being good is not a relational property—that is, not a property whose possession requires that there is a relation of some sort between the things that are good and particular people—it follows that all it can mean to say that being happy is good for *A* is that *A*’s possession of the property of being happy is good in an unqualified sense. And since there can be no ad hoc restrictions on which of the things that are good in an unqualified sense creates obligations, it follows that, if indeed *A*’s being happy is good, then everyone—*A*, me, you, and everyone else as well—has an obligation to do what they can to make *A* happy. It therefore cannot be, as ethical egoism has it, both that *A*’s being happy is good and yet that you and I and everyone else have no obligation to do what we can to make *A* happy.

II. BROAD’S RESPONSE TO MOORE

As I said, Moore’s official charge is that ethical egoism is self-contradictory. For reasons that are no doubt already plain, however, that charge cannot be sustained. Broad puts the point this way:

Suppose . . . that *A* is an ethical egoist. He can admit that, if a certain experience or disposition of his is good, a precisely similar experience or disposition of *B*’s will be also and equally good. But he will assert that it is not his duty to produce good experiences and dispositions as such, without regard to the question of who will have them. *A* has an obligation to produce good experiences and dispositions in himself, and no such direct obligation to produce them in *B* or in anyone else. *A* recognizes that *B* has no such direct obligation to produce them in *A* or in anyone else. This doctrine does not contradict itself in any way. . . . What it does contradict is Sidgwick’s second axiom about goodness and our

obligations in respect of promoting it. This is stated as follows . . .
 “as a rational being I am bound to aim at good generally—so far
 as it is attainable by my efforts—not at a particular part of it.” . . .
 But this does not make ethical egoism *self*contradictory; and, un-
 less Sidgwick’s axiom be self-evidently true, the inconsistency of
 ethical egoism with it does not refute that doctrine.⁸

In other words, since ethical egoism can be thought of as rejecting one of the premises of Moore’s argument against ethical egoism—the premise that if *A*’s being happy is good in an unqualified sense, then everyone has an obligation to do what they can to make *A* happy—it follows that that doctrine can hardly be charged with being *self*contradictory. Broad is surely right about this.

The problem, though, is that Broad seems to think that having in this way saved ethical egoism from the charge of inconsistency, he has saved it simpliciter. Yet this is very difficult to believe. Indeed, given the overwhelming plausibility of supposing that there can be no ad hoc restrictions on which of the things that are good creates obligations—as I see it, this is what provides the rationale for both the second premise of Moore’s argument, and for Sidgwick’s second axiom—it seems to me that, in the light of Broad’s attempt to salvage the doctrine, we should conclude that Moore does indeed refute ethical egoism, notwithstanding the fact that it doesn’t show the doctrine to be self-contradictory. Or rather, we should conclude that he refutes the doctrine, providing, we assume, the truth of the first premise of his argument, the premise that if being *F* is good for some person, *A*, then this entails that *A*’s possession of the property of being *F* is good in an unqualified sense.

III. GENERALIZING MOORE’S ARGUMENT

It should now be clear that Moore’s argument does indeed readily generalize. Consider commonsense morality’s claim that I have a special obligation to provide benefits to my lover, an obligation not shared by others such as you; that you have a special obligation to provide benefits to your lover, an obligation not shared by others such as me; and so on.

If Moore’s argument is successful, then the benefits that each of our lovers receive cannot have a special property of goodness that belongs to me or goodness that belongs to you, for there is no such property. The benefits that they receive can therefore only be good in an unqualified sense. However, if this is right, then it follows that, in supposing that there are such special obligations, commonsense morality assumes that what each of us is obliged to do is a function of only some, not all, of the things that are good. Since that puts an ad hoc

8. *Ibid.*, pp. 45–46.

restriction on which of the good things creates obligations in each of us, it follows that commonsense morality must be mistaken.

To the extent that we find this conclusion difficult to believe, the question we must ask ourselves is where the argument for it goes wrong. We know that the second premise of Moore's argument is unassailable. That only leaves the first. Should we agree that if being *F* is good for some person, *A*, then this entails that *A*'s possession of the property of being *F* is good in an unqualified sense?

IV. AN ALTERNATIVE RESPONSE TO MOORE

As Moore sees things, being good is not a relational property, that is, a property whose possession presupposes the existence of a relation between the things that are good and particular people. Goodness is rather a simple—by which I take it he means, *inter alia*, nonrelational—nonnatural property possessed by the things that are good. But suppose, just for the sake of argument, that being good did presuppose the existence of some such relation, and hence that goodness itself did, in this sense, “belong” to particular people. That would allow us to reject the first premise of Moore's argument.

For example, suppose that the goodness possessed by my being happy presupposed the existence of some sort of relation between my being happy and me, that the goodness possessed by your being happy presupposed the existence of some sort of relation between your being happy and you, and so on. In that case, we would not be forced to admit that the truth of the claim that my being happy is good for me follows from

$$\forall x (x\text{'s being happy is good}).$$

This is the view that Moore accepts, and which he thinks the ethical egoist must accept, too. Instead, we could suppose that it follows from

$$\forall x (x\text{'s being happy is good}_x),$$

where the subscript on the ‘good’ serves simply to make it clear to whom the goodness “belongs.”⁹ For the moment let's not ask what precisely that subscript means. I will eventually make a suggestion about that (Sec. VI). Let's instead pursue the consequences of the fact that ‘good’ is so subscripted, supposing it to be so.

If goodness were a relational property of the sort envisaged, and if the claim just made were the only general evaluative truth, then all of the following claims would follow: my being happy is good_{me}, your

9. Note that this is plainly a distinct kind of relativity from that proposed by Lloyd Humberstone in his “Two Sorts of Agent-Relativity,” *Philosophical Quarterly* 41 (1991): 144–66.

being happy is not good_{me}, your being happy is good_{you}, my being happy is not good_{you}, and so on. But, if this were right, then we would have a way of making sense of what the ethical egoist was trying to get at in claiming, misleadingly, that each person's being happy is the sole good for them, a way that wouldn't commit the ethical egoist to the absurd conclusion that my being happy is the only thing that is good and that your being happy is also the only thing that is good.¹⁰ For what this would really amount to is the suggestion that each person's being happy is the sole good_{that person}. If goodness were a relational property of the sort envisaged, then there would be nothing absurd about that. It would simply amount to the familiar view that, as we would put it nowadays, being happy is a relative value, rather than a neutral value.

In other words, what would then follow from the fact that being *F* is good for *A* is not, as Moore would have it, that *A*'s possession of the property of being *F* is good in an unqualified sense but, rather, that *A*'s possession of the property of being *F* is good_A.¹¹ If we were then to augment the Moorean conception of obligation in the light of this relative conception of goodness as follows:

$\forall x$ (x ought to Φ in certain circumstances C iff Φ ing is that action, of the actions that x can perform in C , that produces the most good _{x}),

then we would have the resources with which to formulate an alternative account of ethical egoism, an account that is not vulnerable to objections of the kind Moore puts forward in his argument against ethical egoism.

According to this alternative account, ethical egoism comprises two claims. First, it comprises a substantive claim to the effect that something of the appropriate kind is a relative good. We have been assuming that the ethical egoist is a hedonist so, given that assumption, this first claim amounts to the suggestion that being happy is such a good.

$\forall x$ (x 's being happy is good _{x}).

10. As Moore put it in the passage quoted above: "What Egoism holds, therefore, is that each man's being happy is the sole good—that a number of different things are each of them the only good thing there is—an absolute contradiction!" Moore, *Principia*, p. 99.

11. Note that I have claimed that being *F* is good for *A* only if *A*'s being *F* is good_A, not that being *F* is good for *A* if and only if *A*'s being *F* is good_A. This is because it seems to me that, at least according to the common understanding of the expression "good for" that is at issue here, *F*'s being good for *A* requires, additionally, that *A*'s being *F* in some way contributes to *A*'s welfare. Yet, for all that I have said so far, *A*'s being *F* might well be good_A without in any way contributing to *A*'s welfare. *A*'s being fully informed might be like this, for example. It might be good_A that *A* is fully informed, even though her being fully informed contributes only negatively to her welfare. In that case *A*'s being fully informed would be good_A but not good for *A*.

Second, it comprises the Moorean conception of obligation as augmented in the light of the relational conception of goodness.

$\forall x$ (x ought to Φ in certain circumstances C iff Φ ing is that action, of the actions that x can perform in C , that produces the most good _{x}).

Putting these two claims together we get

$\forall x$ (x ought to Φ in certain circumstances C iff Φ ing is that action, of the actions that x can perform in C , that makes x happiest),

and this is, of course, just equivalent to the hedonistic version of ethical egoism.

As I said, it seems to me that this version of ethical egoism is not vulnerable to objections of the kind that Moore puts forward in his argument against ethical egoism. This is not, of course, to say that the doctrine so formulated is true. It is, however, to say that it is at least coherent.

Moreover, and much more important, it seems to me that this version of ethical egoism provides us with the logical resources with which to formulate alternative ethical theories, theories such as commonsense morality, which hold that there are both neutral and relative values, and which hold that our obligations are a function of the different weights that these values have in a given choice situation. In the remainder of this article I will defend these claims by considering a range of objections.

V. FIRST OBJECTION: "THERE IS NO WAY OF BALANCING NEUTRAL AND RELATIVE VALUE"

It might be objected that even if I am right that this version of ethical egoism is coherent, that doesn't show that we have the logical resources with which to formulate alternative ethical theories, such as commonsense morality, which purport to balance neutral and relative value against each other. Indeed, it might be thought that the suggested interpretation of ethical egoism makes for an insuperable difficulty in making sense of the idea of such a balancing. Consider how the formulation might go.

In so far as we think that (say) the happiness of our lovers has relative value, we might suppose that we are committed to

$\forall x$ (x 's lover's being happy is good _{x}),

and this, given the augmented Moorean account of obligation together with the simplifying assumption that there are no other relative values, entails

$\forall x$ (x ought to Φ in certain circumstances C iff Φ ing is that action,

of the actions that x can perform in C , that makes x 's lover happiest).

But how are we to represent the idea that the benefit we give to a stranger has neutral value? It might be thought that the Moorean was right about this. In other words, it might be thought that we are also committed to

$\forall x$ (the maximal aggregate comprised inter alia of y 's being happy is good),

which, given the original Moorean account of obligation together with the simplifying assumption that there are no other neutral values, entails

$\forall x \forall y$ (x ought to Φ in certain circumstances C iff Φ ing is that action, of the actions that x can perform in C , that makes the maximal aggregate of people comprised, inter alia, of y happiest).

But, if this is right, then the upshot is that we have failed to provide an account of relative goods such that relative goods and neutral goods can be balanced against each other. What we have formulated is instead a doctrine according to which neutral goods and relative goods are radically incommensurable. Moreover, radical incommensurability of relative and neutral goods entails that the obligations to which they give rise leave us with an irresolvable dilemma.

In order to see that this is so, consider a version of the example described at the outset. Suppose that I face a choice between giving a slightly smaller benefit to my lover by giving her \$20, or a slightly greater benefit to a complete stranger by giving him \$20. Assuming that the consequences of my actions are otherwise identical, let's suppose that commonsense morality tells us that I have an obligation to give the smaller benefit to my lover—that is, to act on the relative value at stake, rather than the neutral value—because the gain in relative value outweighs the loss in neutral value. The trouble is that, on the suggested interpretation of commonsense morality, this is not what I am obliged to do. Instead, on the suggested interpretation, I have inconsistent obligations: an obligation to give the \$20 to the stranger and an obligation to give that same \$20 to my lover. Moreover, there is no way to remove this inconsistency because there is no way to weigh the gain in goodness_{me} that I produce by giving the \$20 to my lover, thereby making her as happy I can, against the loss in good (unsubscripted) that I produce by failing to give the \$20 to the stranger, thereby failing to make him as happy as I can. Goodness (unsubscripted) and goodness_{me} are simply different properties, so there is nothing they have in common that would allow us to compare them. All we can say is that I produce a little bit more goodness_{me} and less goodness (unsubscripted) by giving my lover

\$20, and a little bit less goodness_{me} and more goodness (unsubscripted) by giving the stranger \$20.

How are we to reply to this objection? It might be thought that the best we can do is to look at the obligations themselves and see whether one or another of them is more stringent. Perhaps obligations arising from neutral goods are lexically prior to those arising from relative goods. Or perhaps the reverse is true. Or perhaps one kind of obligation is prior in circumstances of one kind, and the other in circumstances of a different kind. But while some might be attracted to this sort of reply, I am not. The problems are twofold.

First, once we buy into a Moorean conception of obligation, I see no way of analyzing the stringency of an obligation except by way of considering the amount of good that acting on that obligation will produce. The incommensurability of neutral and relative goods thus seems to make it impossible to suppose that one, rather than the other, obligation is more stringent. I therefore see no way of making a reply of this sort fully intelligible.¹²

But there is another and more fundamental problem as well, as it seems to me that this sort of reply concedes far too much. The proper reply to the objection lies in a rejection of the suggested interpretation of commonsense morality. The implicit idea behind the suggested interpretation is that, just as we need to posit goodness_{me}, goodness_{you}, and so on, in order to account for the existence of relative value, so we need to posit goodness (unsubscripted) in order to account for the existence of neutral values. Thus, on the suggested interpretation of commonsense morality, the property of being good comes in two kinds—there is unsubscripted goodness and subscripted goodness—and the distinction between neutral and relative goods is made by distinguishing between the kind of goodness that the properties in virtue of which things are good possess.

There is, however, an alternative interpretation of commonsense

12. This objection seems to me to plague even those analyses of obligation, such as the following, that are structurally like the “hybrid view” Samuel Scheffler presents in *The Rejection of Consequentialism* (Oxford: Clarendon, 1982), pp. 19–21: “ $\forall x$ (x ought to Φ in certain circumstances C iff either Φ -ing is that action, of the actions that x can perform in C, that produces the most good, or, provided that the loss of good isn’t too great as compared with the gain in good, Φ -ing is that action, of the actions that x can perform in C, that produces the most good_x.)” For what could it mean for a loss of good to be “too great,” as compared with a gain in good_x, except that the gain and loss of the relevant amounts of good, and good, results in a situation that is worse? The crucial point is then that ‘worse’ can only mean either ‘worse (unsubscripted)’ or ‘worse_x’, in which case no progress has been made. (Note that I am not suggesting that this is an objection to Scheffler’s own view. My suspicion is that the alternative account of the distinction between neutral and relative value that I go on to describe in the text is one which he could happily accept.)

morality, an interpretation according to which there is only one kind of goodness—there is only subscripted goodness—and the distinction between neutral and relative goods is made by distinguishing between the properties in virtue of which things are good. According to this alternative interpretation, insofar as commonsense morality holds that the happiness of our lovers has relative value, we must suppose that it is committed to

$\forall x$ (x 's lover's being happy is good _{x}).

But then, insofar as it holds that the happiness of strangers has neutral value, we must suppose that it is committed to

$\forall x\forall y$ (the maximal aggregate comprised inter alia of y 's being happy is good _{x}).

Of course, the subscript in the claim just made may look rather idle. But in fact it serves to signal the important fact that goodness is a relational property, a property whose possession presupposes the existence of a relation between the properties in virtue of which things are good, no matter which properties they are, and particular people.

According to commonsense morality, so interpreted, it follows that the difference between neutral goods and relative goods lies not in the kind of goodness that good things possess. Goodness is uniform. Rather, the difference lies in the nature of the properties in virtue of which things are good, that is, the properties that instantiate the property of being good. More specifically, neutral goods are those properties in virtue of which things are good whose characterization requires no mention of the subject to whom the property of being good is a relation. Happiness, no matter who experiences it, is perhaps an example of such a neutral good. Relative goods, by contrast, are those properties in virtue of which things are good whose characterization requires that that subject be explicitly mentioned. My happiness is perhaps an example of such a relative good.

The main advantage of distinguishing between neutral and relative goods in this sort of way is theoretical, as it avoids the ontological extravagance of positing two kinds of goodness: relational goodness and goodness in an unqualified sense, that is, nonrelational goodness. Goodness is a relational property, and that's that. The additional advantage, however, is that distinguishing between neutral and relative goods in this way allows us to answer a question that is crucial to commonsense morality, a question that the earlier formulation of the distinction made impossible to answer. Consider once again the example.

Since my giving \$20 to a stranger produces a certain benefit to the stranger, and since

$\forall x \forall y$ (the maximal aggregate comprised inter alia of y 's being happy is good _{x}),

it follows that the stranger's having that level of benefit has a certain amount of goodness_{me}. Moreover, since my giving \$20 to my lover produces a certain benefit to my lover, and since

$\forall x$ (x 's lover's being happy is good _{x}),

it follows that my lover's having that level of benefit has a certain amount of goodness_{me} as well. The former is a neutral value, because in characterizing what it is about it that makes it valuable, I needn't be mentioned, whereas the latter is a relative value, because in characterizing what it is about it that makes it valuable, I must be mentioned.

Combining each of these claims with the augmented Moorean account of obligation we get conflicting obligations, just as before. But this time we can see how that conflict can be resolved. It makes perfect sense to ask whether there is more goodness_{me} to be gained by my lover's having the smaller benefit, or more goodness_{me} to be gained by the stranger's having the larger benefit. In this way it follows that we can weigh the neutral and the relative goods against each other, just as commonsense morality requires. And since, by assumption in this version of the example, there is more goodness_{me} in my lover's having the smaller benefit than there is in the stranger's having the larger benefit, we can put this conclusion together with the Moorean conception of obligation, as augmented in the light of the relative conception of goodness:

$\forall x$ (x ought to Φ in certain circumstances C iff Φ ing is that action, of the actions that x can perform in C , that produces the most good _{x}),

in order to generate an obligation. What I have an obligation to do is thus to provide my lover with the smaller benefit, just as commonsense morality suggests.

So far I have been concerned simply to formulate an alternative, coherent, interpretation of a theory which holds, as commonsense morality does, that neutral and relative values can be weighed against each other. According to this interpretation the distinction between relative and neutral goods is a distinction among the things that are good, not a distinction between the kind of goodness that the things that are good possess. It is, however, important to realize that that alternative isn't just coherent, but that it is also, for this very reason, far more plausible than the interpretation that requires us to suppose that the distinction is a distinction in the kind of goodness possessed by the things that are good. Moreover, this in turn suggests that we have the makings of an argument against that interpretation.

For one thing, having posited the existence of goodness_{me}, goodness_{you}, and the like, in order to account for the existence of relative values, it is, as we have already said, ontologically extravagant to posit the existence of goodness (unsubscripted) in order to account for the existence of neutral values when, as we have just seen, we can account for the existence of neutral values without positing the existence of goodness (unsubscripted). For another, it is simply perverse to posit a distinction between relative and neutral goods that makes it impossible for relative and neutral goods to be weighed against each other in the way we commonsensically suppose they are. The very fact that they cannot weigh against each other in the way we commonsensically suppose they do therefore counts against supposing that the distinction is to be made in that way.

Having said that, let me emphasize that I have no stake in how we use the terms 'relative' and 'neutral'. These terms are, it seems to me, philosophers' terms of art, and can therefore be used in whatever way we like. If someone insisted that we use the terms 'relative' and 'neutral' to pick out the views that goodness is a relational property and a non-relational property, respectively, then I have no real objection to that. All that I would insist upon is that that use of the terms 'relative' and 'neutral' will be of no help in illuminating the debate between consequentialists, commonsense moral theorists, and ethical egoists. For, as we have seen, in order to illuminate that debate we need terms that distinguish among the views all of which hold in common that goodness is a relational property. Specifically, we need terms that distinguish among the properties in virtue of which things are good. On the one hand, there are properties in virtue of which things are good whose characterization, as good-makers, requires mention of the person to whom goodness is a relation. On the other hand, there are properties whose characterization as good-makers requires no such mention.

VI. SECOND OBJECTION: "THE DEMAND FOR AN INTERPRETATION OF THE SUBSCRIPT CANNOT BE MET"

So far we have been going along with the assumption that the property of being good itself belongs to particular people, in the sense of presupposing the existence of a relation that the things that are good stand in to those people. More specifically, we have been supposing that each person's being happy might be good in the following sense:

$\forall x$ (x 's being happy is good _{x}).

However, we haven't yet attempted to give an interpretation of that subscript. The second objection takes issue with this. The objection is that the demand that we give a plausible interpretation of that subscript cannot be met.

Note that this second objection is, in effect, a replay of the argument for the first premise of Moore's argument against ethical egoism. According to Moore, goodness is not a relational property, relating things that are good to particular people, but is rather metaphysically simple: it is a simple, nonnatural property possessed by things that are good. Of course, these days few people believe that Moore is right that goodness is such a simple, nonnatural property. But that does nothing to show, more positively, that goodness is a relational property of the kind required. The task then, is to make a plausible case for that conclusion.

There may well be a variety of ways in which we could conceive of the property of being good as a relational property of the required kind.¹³ However, my own view is that the best way of doing so is by giving a detailed statement and defense of a particular version of the dispositional theory of value.¹⁴ Doing this would, however, take us too far afield in the present context, so I will simply give a rough sketch of the version of the dispositional theory that I myself favor and show why it does the job required. Since I have defended the theory at length elsewhere, what I say here will hopefully remind people of how, as I see things, a fuller defense might go.¹⁵

According to the version of the dispositional theory of value that I myself favor, when a subject judges p 's being the case in certain circumstances C to be good, what she is doing is making a judgment about what she would desire to be the case in C if she had a psychology that eludes all forms of rational criticism. In order to fully spell out the dispositional theory we therefore need an account of the nature of a psychology that has this feature. My suggestion in this regard is, very roughly speaking, that a psychology that eludes all forms of rational criticism is a psychology that comprises a desire set that is maximally informed and coherent and unified. When a subject judges p 's being the case in C to be good, what she is judging is therefore, I suggest, that she would desire that p in C if she had a desire set that is maximally informed and coherent and unified.

If this is right, however, then, as is perhaps already clear, it turns out that 'good' is indeed subscripted in just the way required. For when I judge p 's being the case in C to be good, I am judging that p 's being the case in C has a certain relational property, specifically, the property

13. In his contribution to this issue of *Ethics*, Thomas Hurka provides an alternative conception which he finds in Sidgwick.

14. See Michael Smith's (pp. 89–111), David Lewis's (pp. 113–37), and Mark Johnston's (pp. 139–74) contributions to the "Dispositional Theories of Value" symposium in *Proceedings of the Aristotelian Society*, suppl., 63 (1989): 89–174.

15. See, especially, Michael Smith, *The Moral Problem* (Oxford: Blackwell, 1994), and "In Defense of *The Moral Problem*: A Reply to Brink, Copp, and Sayre-McCord," *Ethics* 108 (1997): 84–119.

of being what I would desire if I had a set of desires that was maximally informed and coherent and unified; and when you judge p 's being the case in C to be good, you are judging that p 's being the case in C has a different relational property, the property of being what you would desire if you had a set of desires that was maximally informed and coherent and unified; and so on. In other words, I am really judging p 's being the case in C to be good_{me} and you are really judging p 's being the case to be good_{you}. Our judgments are thus appropriately relational.

Now it might be thought that there is a devastating objection to the dispositional theory and that the objection is a direct corollary of supposing that goodness is a relational property of the kind described, that is, a direct corollary of the apparent egocentricity of such judgments. "After all," the objection goes, "if what you have just said were true then we could never agree or disagree with each other about what is good: I would forever be talking about what I would desire, you would forever be talking about what you would desire, and so on. Yet that is surely quite incredible. Agreement and disagreement about what is good is possible." However, contrary to the objection, it seems to me that the version of the dispositional theory that I have just sketched doesn't entail the impossibility of agreement and disagreement about what is good.

Suppose that I judge that my being happy is good_{me} and that you judge that your being happy is not good_{you}. In other words, I judge that I would desire that I be happy if I had a set of desires that was maximally informed and coherent and unified, and you judge that you would not desire that you be happy if you had a set of desires that was maximally informed and coherent and unified. Do we agree or disagree? In order to answer this question it seems to me that we must first know how each of us would universalize our respective judgments. It is, after all, a conceptual truth that judgments about the good are universalizable.¹⁶ In this respect, there would seem to be at least two possibilities.

One possibility is that, when I universalize, I find that I am committed to the value theory of ethical egoism:

$$\forall x (x\text{'s being happy is good}_x).$$

That is, in terms of the dispositional theory, I am committed to

$$\forall x (\text{if } x \text{ had a set of desires that was maximally informed and coherent and unified, then } x \text{ would desire that } x \text{ be happy}).$$

Alternatively, when I universalize, I might find that I am committed to the value theory of utilitarianism:

16. This is a recurring theme in R. M. Hare's *The Language of Morals* (Oxford: Oxford University Press, 1952), *Freedom and Reason* (Oxford: Oxford University Press, 1963), and *Moral Thinking* (Oxford: Oxford University Press, 1981).

$\forall x \forall y$ (the maximal aggregate comprised inter alia of y 's being happy is good _{x}).

That is, in terms of the dispositional theory, I am committed to

$\forall x$ (if x had a set of desires that was maximally informed and coherent and unified, then x would desire that $\forall y$ [the maximal aggregate comprised inter alia of y be happy]).

Moreover, when you universalize, it seems that you would find that you are committed to a denial of one or the other of these two possibilities.

For present purposes it doesn't matter to which of these we find ourselves committed. The crucial point is simply that when you and I each universalize our judgments, we would therefore find ourselves committed either to accepting and rejecting the same claims, or to accepting and rejecting different claims, and this, in turn, underwrites the possibility of agreement and disagreement.

For example, I might judge that my being happy is good_{me} (i.e., that I would want that I be happy if I had a maximally informed and coherent and unified desire set), because I judge that $\forall x$ (x 's being happy is good _{x}) (i.e., that $\forall x$ [x would want that x be happy if x had a maximally informed and coherent and unified desire set]), whereas you judge your being happy is not good_{you} (i.e., that you would not want that you be happy if you had a maximally informed and coherent and unified desire set) because you reject this claim. If so, then it turns out that we disagree, notwithstanding the apparent egocentricity of our original judgments.

Alternatively, we might find that we make our original judgments because we accept different claims. Perhaps I make mine because I judge that $\forall x$ (x 's being happy is good _{x}) (i.e., that $\forall x$ [x would want that x be happy if x had a maximally informed and coherent and unified desire set]), whereas you make yours because you deny that $\forall x \forall y$ (the maximal aggregate of people comprising inter alia y 's being happy is good _{x}) (i.e., deny that $\forall x$ [x would want that $\forall y$ (the maximal aggregate comprising inter alia y be happy) if x had a maximally informed and coherent and unified desire set]). In that case we might well find that we agree with each other, again notwithstanding the apparent egocentricity of our original judgments, for we might find that we are each disposed to make the same universal judgments.

To sum up, the second objection is that it is impossible to give a plausible interpretation of goodness as a relational property, a property whose possession presupposes the existence of a relation between the things that are good and particular people. But, as I have argued, this is not so. The dispositional theory provides us with an extremely plausible interpretation of this idea. Moreover, though it might initially seem

like supposing that goodness is such a relational property would make it impossible for there to be agreement or disagreement about what is good, we have seen this too is a mistake. Notwithstanding the apparent egocentricity of judgments about the good, because they are universalizable we can see that both agreement and disagreement are possible.

VII. THIRD OBJECTION: "THE POSITED CONNECTION BETWEEN GOODNESS AND OBLIGATION IS AD HOC"

The third objection is that when we combine the proposed augmented Moorean conception of obligation

$$\forall x (x \text{ ought to } \Phi \text{ in certain circumstances } C \text{ iff } \Phi \text{ing is that action, of the actions that } x \text{ can perform in } C, \text{ that produces the most good}_x)$$

with the claim that there are relative goods, we become guilty of putting ad hoc restrictions on which of the things that are good give rise to obligations. As is perhaps evident, this third objection recalls the argument for the second premise of Moore's argument against ethical egoism.

In order to appreciate the full force of this objection, imagine combining the augmented Moorean conception of obligation with the claim that being happy is a relative good:

$$\forall x (x\text{'s being happy is good}_x).$$

If this were true, then I would have to accept all of the following: that my being happy is good_{me}; that your being happy is good_{you}; that (say) Φ ing, and not Ψ ing, is the action of those I can perform which will produce the most good_{me}; and that (say) Ψ ing, not Φ ing, is the action of those I can perform which will produce the most good_{you}; and so on. In other words, I would have to accept a whole variety of claims about which things are relative goods and about the actions that I can perform that will produce various amounts of these relative goods.

However, and crucially, I would also have to accept that only certain of these claims, those that are specifically indexed to me, generate my obligations. But why should this be so? Why shouldn't we suppose instead that (say)

$$\forall x \forall y (x \text{ ought to } \Phi \text{ in certain circumstances } C \text{ iff } \Phi \text{ing is that action, of the actions that } x \text{ can perform in } C, \text{ that produces the most good}_y).$$

In other words, why don't goods_{you} oblige me just as much as goods_{me}? Isn't it simply ad hoc to suppose, as we must if we accept the proposed augmented Moorean conception of obligation, that only goods_{me} oblige

me? Isn't this a violation of Sidgwick's second axiom requiring us to aim at the whole of the good, not simply a part of it?

The reply to these questions comes in two parts. Let's begin by restricting ourselves to neutral goods, that is, to properties in virtue of which things are good whose characterization does not require that we mention the subscripted person. Suppose, for example, that

$\forall x \forall y$ (the maximal aggregate comprised inter alia of y 's being happy is good _{x}).

From this it follows immediately that

$\forall y$ (the maximal aggregate comprised inter alia of y 's being happy is good_{me} and good_{you}).

But in that case, at least when neutral goods are at stake, it follows that there is no difference between Φ ing's being that action, of the actions that x can perform in C , that produces the most good _{x} and Φ ing's being that action, of the actions that x can perform in C , that produces the most good _{y} . No matter how we define our obligations, when the goods are neutral, goods_{you} oblige me just as much as goods_{me}. There is no issue to be engaged.

When we consider relative goods, however, there is an issue to be engaged. For, when the things that are good _{x} are relative goods, there is all the difference in the world between Φ ing's being that action, of the actions that x can perform in C , that produces the most good _{x} and Φ ing's being that action, of the actions that x can perform in C , that produces the most good _{y} . If, for example, we suppose that

$\forall x$ (x 's being happy is good _{x}),

then the difference is that between my being obliged to do what I can to make myself as happy as I can be while having no such obligation to make you happy, and my being obliged to do what I can to make you as happy as you can be as well. The question that we must answer is thus why, when it comes to relative goods, it is so important that my obligations should be defined in terms of the goods_{me}, rather than the goods_{each person}. Isn't this ad hoc? I do not think that it is, and the reason why lies in a feature of the dispositional theory that I have explained at greater length elsewhere. Let me first describe that feature, and then explain how it provides us with the needed explanation.

One of the main attractions of the dispositional theory, as I see things, is that it allows us to explain why the doctrine known as "judgment internalism" is true, that is, why it is true that, absent practical irrationality, a subject who judges p 's being the case to be good desires

that p be the case.¹⁷ The argument for this conclusion is what has come to be known as the “Incoherence Argument.”¹⁸ The Incoherence Argument runs roughly as follows.

Let’s agree that when someone judges p ’s being the case to be good_{her} what she believes is that she would desire that p if she had a maximally informed and coherent and unified set of desires. But now consider the following pairs of psychological states that that subject might have: the belief that she would desire that p if she had a maximally informed and coherent and unified set of desires together with a desire that p ; that same belief together with indifference toward p ; and that same belief together with an aversion toward p . Which of these pairs of psychological states is more coherent? It seems that the answer to this question is plain. The first of these pairs is much more coherent than the second or the third.

If this is right, however, then it follows immediately that if that subject is practically rational, in the sense of having and exercising a capacity to have the psychological states that coherence demands of her, then—at least abstracting away from such other dynamic changes in her beliefs as might occur for evidential reasons—she will end up having a desire that matches her belief about what she would want if she had a maximally informed and coherent and unified desire set.¹⁹ In other words, in the particular case under discussion, she will end up desiring p .

Judgment internalism is thus true of the judgments that a particular subject makes about which things are good_{her} because, absent practical irrationality—that is, absent a failure to have or exercise the capacity to have a coherent psychology—a subject who believes that she would desire that p if she had a maximally informed and coherent and unified set of desires desires that p . Now let’s consider the judgments a subject makes about things that are good_{someone else} in order to see whether that same argument shows that judgment internalism is true of these judgments as well.

As we have already seen, when neutral values are at stake, everything a subject judges to be good_{someone else} is something that she also judges to be good_{her}. Judgment internalism is thus plainly true of the judgments

17. For more on judgment internalism, see Stephen Darwall, *Impartial Reason* (Ithaca, N.Y.: Cornell University Press, 1983); and, especially, Christine Korsgaard, “Skepticism about Practical Reason,” *Journal of Philosophy* 83 (1986): 5–25.

18. The first explicit formulation of this argument is in my “In Defence of *The Moral Problem*.” The argument is discussed in Russ Schafer-Landau, “Moral Judgement and Normative Reasons,” *Analysis* 59 (1999): 33–40; and my “The Incoherence Argument: Reply to Schafer-Landau,” in *Analysis* 61 (2001): 254–66.

19. Nomy Arpaly, “On Acting Rationally against One’s Best Judgment,” *Ethics* 110 (2000): 488–513, discusses the dynamic changes abstracted away from.

a subject makes about things that are good_{someone else}, but it is true of them precisely because it is true of the judgments she makes about things that are good_{her}. But what about those judgments a subject makes about things that are good_{someone else} where relative values are at stake?

Let's suppose that a subject believes that his own happiness is good_{him} because he believes that he would desire that he be happy in certain circumstances *C* if he had a maximally informed and coherent and unified desire set, and hence believes that someone else's happiness is good_{that other person} because she would desire that she be happy in certain circumstances *C* if she had a maximally informed and coherent and unified desire set: in other words, he believes that $\forall x$ (x would desire that x be happy if x had a maximally informed and coherent and unified desire set). Now consider the following pairs of psychological states that the original subject might have: his belief that that other subject would desire that she be happy in *C* if she had a maximally informed and coherent and unified desire set together with a desire, on his own behalf, that she be happy in *C*; that same belief together with indifference, on his behalf, toward her being happy in *C*; and that same belief together with an aversion, on his behalf, toward her being happy in *C*. Is the first of these pairs of psychological states more coherent than the others?

The answer would seem to be that it plainly isn't more coherent than the others. There is no incoherence at all in a subject's believing that someone else would desire that she be happy in *C*, if she had a maximally informed and coherent and unified desire set, without desiring himself that she be happy in *C*. This is because, given what the subject believes about the other subject, it follows that he must believe that he himself would desire that he, not she, be happy in *C* if he had a maximally informed and coherent and unified desire set. But in that case, what coherence plainly requires of him, for much the reasons given in the original presentation of the Incoherence Argument, is that he desires that he be happy in *C*, not that he desires that she be happy in *C*.

We are now in a position to give the second part of our reply to the third objection. The crucial point is that judgment internalism is true of the "good" claims I accept that are indexed to me but not true of the "good" claims I accept that are indexed to other people, because there is no incoherence at all involved in my having beliefs about things that are good_{other people}, and yet having no desire at all that those goods_{other people} obtain (or, more precisely, there is no incoherence so long as we cannot infer from the fact that the thing that is good_{other people} that it is also good_{me}, as is the case with goods that are neutral). There is, however, plainly incoherence involved in my having beliefs about things that are good_{me} and yet having no desire that those goods_{me} obtain. This, in turn,

explains why we have to analyze each person's obligations in terms of the things that that person can do that produce the most good_{that person} in the way suggested by the augmented Moorean conception of obligation.

Obligations are plainly normative—that is to say, the obligations people have make a claim on them—and the claim that they make is that captured by judgment internalism. The reason we do not have obligations to produce such goods_{other people} as we can is, thus, that such goods make no claim upon us (or, more precisely, goods_{other people} that aren't also goods_{us}, as is the case with neutral goods, make no claim on us). Goods_{other people} (or, more precisely, those goods_{other people} that aren't also goods_{us}) make no claim on us because there is no practical irrationality involved in our acknowledging that we can act so as to promote such goods without being correspondingly motivated. But since the idea of an obligation that makes no claim on us is incoherent, it follows that there is nothing ad hoc about the suggestion that our obligations are restricted to the things that we can do that promote goods_{us}.

VIII. CONCLUSION

As I said at the outset, Moore provides us with an attractive analysis of obligation in terms of goodness and badness. On the basis of this analysis, he goes on to argue that ethical egoism is incoherent. His argument, if successful, would readily generalize to show that all ethical doctrines that posit relative values, either alongside or instead of, neutral values, are likewise incoherent. My main aim in this article has been to explain where Moore's argument against ethical egoism goes wrong. My suggestion has been that it goes wrong in presupposing, mistakenly, that goodness and badness are nonrelational properties. This presupposition is mistaken because the best theory of value, the dispositional theory, tells us that goodness and badness are relational properties. Once we correct this mistaken presupposition, and amend the attractive Moorean analysis of obligation, we see not just that ethical egoism is coherent but also what the distinction between neutral and relative values really consists in, and what it is about our obligations that makes them make a claim on us.