

CONSENSUS AND EXCELLENCE OF REASONS

JEREMY RANDEL KOONS

AMERICAN UNIVERSITY OF BEIRUT

ABSTRACT: It is plausible to suppose that the normativity of evaluative (e.g., moral and epistemic) judgments arises out of and is, in some sense, dependent on our actual evaluative practice. At the same time, though, it seems likely that the correctness of evaluative judgments is not merely a matter of what the underlying practice endorses and condemns; denial of this leads one into a rather objectionable form of relativism. In this paper, I will explore a social practice account of normativity according to which normativity is grounded in our actual social practice of evaluation. I will show how this account allows normativity to be dependent on our actual evaluative practice, while allowing the correctness of evaluative judgments to be independent of this practice in important ways, and how the resulting temporal logic of reasons gives us a conception of morality and other sorts of evaluative discourse that is not historically local.

T

here are different accounts of what it is for a perceived normative requirement to be correct. These accounts can be seen as lying on a continuum, from those divorcing correctness completely from what the community actually does to those that completely reduce moral correctness to what the community does. The former group may be called transcendentalists¹; they think that what is in fact, for example, morally correct in no way depends on what the community takes to be morally correct; questions of the correctness of normative judgments (moral, epistemic, semantic, etc.) are completely divorced from the actual practice of the community. Under the heading of “transcendentalist” fall Platonists and other assorted third-realmers.

At the other end of the spectrum we have attributivists, those who think that the correctness of a moral judgment is merely a matter of what a particular

community thinks, or of what a particular person thinks. That is, normative correctness is reducible to some community's practice, or to some person's psychological states. Into this camp fall relativists such as Ruth Benedict, along with subjectivists, various emotivists, and others of their ilk.

It is plausible to suppose that the truth of the matter lies somewhere in the middle of the continuum: that is, the correctness of evaluative judgments (e.g., questions of moral rightness and wrongness, epistemic justification, etc.) is in some sense dependent on our actual evaluative practice. At the same time, though, this dependence is not a tight one: the correctness of evaluative judgments is not merely a matter of what the underlying practice endorses and condemns; denial of this leads one into a rather objectionable form of relativism. The question then becomes, What are the details of this *via media*? How do we sketch the relation of dependence between the correctness of evaluative judgments and our actual evaluative practice in a way that does justice to both of these intuitions?

In this paper, I will explore a social practice account of normativity—an account according to which normativity is not transcendent in origin, but is instead grounded in our actual social practice of evaluation. I will show how this social practice account of normativity does justice to both of the intuitions I have described—that is, how this account allows normativity to be dependent on our actual evaluative practice, while independent of it in important ways—and how the resulting temporal logic of reasons gives us a conception of morality and other sorts of evaluative discourse that is not historically local. In other words, a social practice account of normative judgment can avoid both attributivism and historical localism.

I. TWO CONCEPTIONS OF NORMATIVITY

Traditional theories of normativity divide into two competing camps. A social practice view of normative utterances is most plausible if we view it as arising in the context of these two views, and as responding to difficulties the two traditional views cannot address. The first camp is the attributivist camp. On this conception of normativity, to say that one ought to do x is to say that the doing of x follows from some set of rules, or that the community endorses the performance of x , or that the performance of x follows from one's beliefs (or would after suitable reflection). Thus, the attributivist conception is essentially a descriptive notion of normativity. To say that one ought to do x is to describe something, either that person's belief set, or the practices of her community, or a set of rules, or some such. It is to say that the endorsement of x follows *as a matter of fact* from this practice or set of beliefs or rules.

The other traditional conception of normativity is the transcendental conception. Far from asserting that particular endorsements follow from a community's practice or an individual's belief set, this view claims that

normative claims aim to correspond to community-independent norms. The correctness of a normative judgment is then entirely independent of what anyone in the community thinks or does.

These are only rough sketches of the two positions, but more is not needed, for certain examples clearly confound both conceptions, no matter how these conceptions are ultimately fleshed out. A good example is the introduction of the “Wilt Chamberlin” rules in basketball.

These were a set of rules introduced for the simple, albeit theoretically puzzling, reason that he was too good a player. His ability to swat balls away from the basket led to both the rule against goaltending and the extension of the key. His ability to leap from behind the foul line and stuff the ball led to a revision in the rules regarding foul shots.²

This example brings out the inadequacy of the attributivist conception of norms:

[I]f claiming that an act is permissible in a game is just saying that this entitlement follows from some fixed set of rules definitive of that game, then it would seem we must conclude that a different game was being played post-Wilt from the one played pre-Wilt. But is this a reasonable conclusion? If so, then we must say that every record book is mistaken since they all take statistics before and after Wilt to have been set in the same game. Everyone who talks about basketball would be subject to a crucial ambiguity of which they are totally unaware. Presumably all sorts of legal issues concerning, say, television rights over *basketball* games would have to be reconsidered.³

Nor can it be said that these revisions were uniquely determined by the state of the game pre-Wilt. Not only is it highly unlikely that it occurred to anyone that such a good player might come along, but these particular solutions to the Wilt problem were underdetermined by the practice of basketball pre-Wilt. In other words, “[t]here were many ways the game could have been revised to retain its interest. (The baskets could have been raised, players over 7 feet tall could have been banned, etc.)”⁴

This “normative underdetermination of the future emendations of the rules by the present needs and past practices demonstrates not only the inadequacy of the attributive conception, but that of the transcendental conception as well. There cannot be a single basketball game out there waiting for us in Plato’s heaven since there are several—indeed, probably infinitely many—coherent games we might have opted for.”⁵

Thus, the transcendentalist and the attributivist are making opposite errors. The transcendentalist divorces normativity too far from social practice, claiming that our practice has no bearing on normative truth. The attributivist, on the other hand, makes the connection too tight, so that any change in our actual practice makes it the case that we are simply engaging

in another activity. This suggests that the correct path is a sort of *via media*, on which normativity arises out of our social practice, but according to which the correctness of a normative claim isn't just a matter of whatever I (or we) say it is, or of whatever follows from my belief set, etc.⁶ In what follows, I will demonstrate how a social practice account can avoid both attributivism and historical localism. This is an important result, because it will show that although crude social practice accounts may run afoul of one or both of these problems, social practice accounts of normativity need not be seen as inevitably succumbing to them.

Before continuing, let me make some brief remarks about social practices themselves. I think we all have a rough, intuitive idea of what a social practice is. It is, for example, what anthropologists posit to explain certain regularities in the behavior of a community. Here is one example: “[C]onsider the practice in basketball of not stepping outside the sidelines when one is in possession of the ball. It is clearly useful as part of a systematic account of the ‘dance of the basketball players’ to take them to be committed to following such a rule, even though they do not always follow it, nor does any penalty for non-compliance universally follow.”⁷ We attribute commitments such as these to communities whose behavior we seek to explain. It is the fact that a bit of behavior (staying inside the lines, stopping at stop signs) is caught up in such a web of commitments that makes possible its classification as an *action* (and its performer as an *agent*), as opposed to mere behavior. In interpreting a community in terms of social practices, one takes the members of the community to be committed to, and bound by, certain appropriatenesses. The normativity is only shallow, however, as positing such practices does not involve committing oneself to following them.⁸

One advantage of explaining norm-governed behavior in terms of social practices is that such an explanation allows us to account for normativity without positing spooky non-natural properties or objects; we need only posit social practices, which are sets of commitments to practical appropriatenesses (mainly implicit), commitments we posit to explain the behavior of a community. Acting according to one of these appropriatenesses is acting for a *reason*. And so reasons are, on this account, products of social practices and the appropriatenesses implicit in these practices.

Despite the advantages of a social practice account of normativity, it seems as though there is a significant disadvantage, however: it seems as though this account leads us into a rather nasty version of relativism. After all, if it is our community's practice that fixes the correctness of an action, then it seems straightforwardly contradictory to question whether an action that accords with our community's practices is correct. We will see, over the course of this paper, how a sophisticated social practice account can evade this worry.

II. MORAL TRUTH

So we have an account of where reasons come from; they come from the practical appropriatenesses implicit in social practices. How do we get from this account of reasons to an account of the truth of a normative claim? In other words, what makes a perceived normative requirement correct? One suggestion lies in Wiggins⁹ and McDowell's¹⁰ discussion of moral truth. Wiggins and McDowell have suggested that we can make sense of moral truth in terms of excellence of reasons. To support this picture, Wiggins draws an analogy with mathematics. He writes that such reasoning could explain the great degree of convergence displayed by people on the claim that $7 + 5 = 12$. This belief (that $7 + 5 = 12$) "resembles an ordinary empirical belief in being *uncontroversially true* . . . but resembles a moral belief in *not being empirical*."¹¹ According to Wiggins, the best explanation for why people believe this is that it can be shown by the calculating rules to be true. "There is nothing else to think but that $7 + 5 = 12$."¹² Thus, even if we cannot posit mathematical facts or entities that stand in a causal relation to us (and are thereby capable of satisfying Harman's explanatory requirement¹³), we can make sense of mathematical truths in terms of excellence of reasons. The reasons for thinking that $7 + 5 = 12$ are so compelling that there is nothing else to think. To restate this in non-metaphorical terms, the best reasons support the conclusion that $7 + 5 = 12$.

Wiggins and McDowell seem to think that their account of truth in terms of excellence of reasons can be extended to every area of discourse, from moral to scientific. Wiggins writes:

A subject matter is objective or relates to objective reality if and only if there are questions about it that admit of answers that are substantially true. It is sufficient for some judgment that p to be substantially true that one could *come to know* that p . One can come to know that p only if one can come to believe that p precisely because p . And one comes to believe p precisely because p only if the best full explanation of one's coming to believe that p requires the giver of the explanation to adduce in his explanation the very fact that p . What follows from this is that his explanation will conform to the following schema: for this, that and the other reasons (here the explainer specifies these), there is really nothing else to think but that p ; so, given the circumstances and given the subject's cognitive capacities and opportunities and given his access to what leaves nothing else to think but that p , no wonder he believes that p . Let us call such an explanation of a belief a *vindictory explanation* of that belief.¹⁴

The view that we can make sense of moral truth in terms of excellence of reasons seems quite plausible. Consider an analogy from chess. A certain

arrangement, *A*, of pieces on the board constitutes checkmate. But there is nothing intrinsic to these pieces of wood on the board that they should, in this particular arrangement, constitute a threat to that particular piece of wood. Rather, *A* constitutes checkmate *in virtue of the rules of chess*. Crudely put, the rules make it the case that *A* constitutes checkmate (indeed, the term “checkmate” has no meaning itself in abstraction from the rules of chess). Morality is in an important sense analogous to games like chess. There’s nothing intrinsic to a particular behavior, *B*, by a featherless biped that makes this behavior evil; rather, it is evil in virtue of the rules of morality. Now, there are those (i.e., particularists) who deny that morality can be codified into a set of rules, so the phrase “rules of morality” will rub these people the wrong way. Let me, then, use the term “reasons” more broadly, to apply to non-codifiable reasons *or* codified rules. This isn’t an abuse of the term “reason” because rules can serve as reasons: for the particularist, judgments are backed by reasons (Why was that action cruel? Because it caused her gratuitous pain). These reasons play the same role for the particularist that rules play for the philosopher who thinks morality can be codified (Why was that action cruel? Because causing gratuitous pain is cruel). That is to say, both reasons and rules justify moral judgments. More important, though, just as the rules of chess make *A* checkmate, the reasons (read: rules or reasons) of morality make *B* cruel; there is nothing intrinsic to *B* that makes it cruel. It is only so in virtue of our moral practice, and the appropriatenesses that constitute this practice.¹⁵

There are reasons to think that the account of truth offered by Wiggins and McDowell is not entirely satisfactory as it stands. First, both seem to think that what this vindicatory explanation must be explaining is convergence in a particular belief. But I don’t think that realism requires the possibility of rational convergence. For example, few (except for a few verificationists) would deny that there is a fact of the matter regarding whether the number of neutrinos in the universe is even or odd. There is, however, not even the slightest possibility of rational convergence on this matter. Consider another example, the venerable brain in a vat. The brain (which we shall call “*X*”) and I will never converge on the claim, “*X* is a brain in a vat.” Even if I were somehow able periodically to hook myself up to the machine and enter the brain’s extended hallucination as a character and argue extensively with the brain, the brain would remain unconvinced, and with considerable justification. It would regard me as I would regard some crank who kept approaching me on the street and arguing that I was a brain in a vat. It seems, then, that what needs to be explained by the vindicatory explanation is not convergence on the belief, but instead individual instances of the belief: *my* belief that *X* is a brain in a vat is best explained by the fact that *X* is, in fact, a brain in a vat. If we can explain how I came to hold *p* in a way that requires *p* to be true, then we can be realists about *p*.

A second consideration shows, however, that even this condition (that an individual might come to believe p precisely because p) is too strong. This fact is best illustrated by the above neutrino example: it might be impossible to come to have a belief about the number of neutrinos in the universe in a way best explained by this being the correct number. Yet there is an objective fact of the matter as to whether the number of neutrinos in the universe is even or odd. Thus, the criterion for realism that Wiggins and McDowell suggest cannot work in all cases. However, I only wish to draw a single insight from their approach: we can make sense of moral truth in terms of excellence of reasons. Let us further examine how this account can be made to work for morality.

III. MORAL TRUTH AND THE TEMPORAL LOGIC OF REASONS

We have seen that on a social practice account, reasons have their origin in the practical appropriatenesses implicit in a social practice. Further, we have seen how Wiggins and McDowell explain moral truth in terms of excellence of reasons. If we combine these two conclusions, the resulting view of normative discourse has a distinct advantage: it allows us to reconcile two theses that are, on their face, difficult to reconcile. The first claim is that morality and epistemology are a creation of humans, that moral and epistemic predicates are dependent on our social practices. The second claim is that the truth of a normative claim (such as “That was immoral”) can often be entirely independent of anything people say or believe about this claim. Let us see how, given this account of normative discourse, normative truths can be dependent on social practice in some sense, yet independent of it in another.

Social dependence (i.e., dependence on a social practice) can be strong or weak. If which performance is correct is merely a matter of which performance most people take to be correct, then that correctness is strongly socially dependent. Consider, for example, the town in Virginia named Buena Vista. Contrary to what anyone with a passing acquaintance with the Spanish language might think, the name of this town is pronounced “Byoona Vissta.” What’s more, this is the correct way to pronounce the name of this town. It would be wrong to pronounce the name of this town as a Spanish-speaker would. Or, to take a more intuitive case, consider the state of Illinois. According to my dictionary, this is a word of French origin. But it would be incorrect to pronounce “Illinois” as a French person would (“Ee-yen-wah,” or some such). In both cases, the pronunciation of the word in question is correct in virtue of what people take the correct pronunciation to be. If Virginians pronounced “Buena Vista” differently than they do, or if Americans pronounced “Illinois” differently than they do, then *that* pronunciation would be the correct one. In short,

if most people in the relevant community take a certain pronunciation to be correct, then that pronunciation *is* correct.¹⁶

There are performances whose correctness does not depend on what most people take to be correct. The truth of some propositions is not a question of what most people believe about these propositions. The paradigmatic cases of such propositions are scientific ones. For example, water would be identical to H₂O even if no one believed that it was. Whether water is identical to H₂O is a socially independent fact. However, it is difficult to understand how social dependence can be anything other than strong social dependence. I wish to argue, though, that morality and epistemology are, as I will call it, “weakly socially dependent.” That is, both are socially dependent but most, or even all, of us could be wrong about the truth of a given moral or epistemological claim. For ease of phrasing, I will discuss this weak social dependence in terms of morality, but it will be easy to apply our conclusions to epistemology and other sorts of normative discourse as well.

An important fact to note, if we are to head off worries of relativism, is that while rules are conventional, they aren’t arbitrary. There are good and bad rules, and we might think of compelling reasons to change some of the rules of the game. To revisit an earlier example, the National Basketball Association used to allow goal-tending. But it became too easy for players such as Wilt Chamberlin to stand under the basket and block shots, so it was decided that goal-tending should be disallowed. It was determined that the ends of the game would be better served by disallowing goal-tending than by allowing it.

Of course, certain elements of basketball practice are arbitrary: the decision to have the point of the game the throwing of a ball through a horizontal hoop was arbitrary. But within the context given by a certain set of arbitrary rules, there are better and worse ways to structure the game. Indeed, one way of looking at the process of creating rules is as a process of codifying reasons. If we make the basket narrower than two basketball-widths, the game will be too hard, and it will be too easy if we make it wider, so we’ll make the basket have a width of two basketballs. Also, if we allow people to carry the ball, that will favor the offense too much, so we’ll require them to dribble. If we allow charging, that will make the game too rough and the incidence of injury too high, so we’ll disallow that. The examples can be multiplied indefinitely; given the minimal framework of arbitrary rules that provide initial structure to the game, the rules of basketball have rationales behind them. The same is true of morality: the decision to drive on the right side of the road (in the U.S.) was arbitrary, but within the framework given by such arbitrary rules, it is clearly immoral to drive on the wrong side of the road, since such activity needlessly endangers others. Thus, certain aspects of social practice are arbitrary, but they form the framework in which objective rules can operate. Indeed, it might

be instructive (even if not strictly accurate) to view morality as an objective function from such arbitrary rules and contingent features of society and the make-up of agents to a set of prescribed and proscribed behaviors. The function is determined by our interests—our interest in not dying provides the connection between the custom of driving on the right side of the road and the moral rule against driving on the left side of the road.¹⁷ Notice, though, that even these arbitrary features can be criticized on moral grounds. As we will see below, we ought to regard no feature of our practice as *de jure* unrevisable.

There is an important disanalogy between morality and games such as basketball. In basketball, all of Wilt Chamberlin's goal-tending before the rule revision was permissible. Since the revision, though, it is a foul. But it would be unfortunate if we had to regard morality in this way. This would have the result that slavery was permissible in the U.S. before the Civil War, and impermissible afterward (think of the Civil War as a big rule revision). But we need not draw this unfortunate conclusion: even if moral truth depends on reasons, this doesn't directly entail that moral truth is historically local. To make my case that basketball and morality are disanalogous in this respect, I must first discuss a feature of reasons: reasons display a sort of "timelessness." Let's look at how the practice of reason-giving works. Suppose we are trying to figure out whether to do *A* (stay home) or *B* (go to the beach), and we think we ought to do *B*. Assume there is a decisive reason *R* (the weather service is saying that a massive storm is moving into the area within the hour) that supports performing action *A*. *R* can be a particularist reason, such as "That caused her gratuitous pain," or a rule such as "In battle, you should never leave your flank unprotected." If someone utters *R*, thereby convincing us that we really ought to do *A*, then we are not justified in thinking, "*B* was the best course of action *until* his utterance made *A* the best course of action." Rather, we are justified in thinking that *A* was best *all along*; it merely took his presentation of *R* to make us realize it. Suppose nobody thinks of *R* at the time, and we go on to do *B*. Even centuries later, someone might truly utter, "They ought to have done *A*," and present *R* to support this conclusion. (It is crucial to distinguish between whether a reason obtained at time *t* and whether agents at time *t* can be blamed for failure to recognize this reason. In other words, we must keep in mind the distinction between "morally wrong" and "blameworthy.") So we can now say both that Galileo's observations provided good reason to think the earth orbits the sun (even though few at the time regarded his observations as such); and we can also give good reasons for thinking that the Church's treatment of Galileo was unjustified. In both cases, we are applying reasons "retroactively," scientific reasons in the one case, and moral and epistemic reasons in the other. But in an important sense, we are not really applying them retroactively; rather, we are presenting reasons which were obtained

at the time of Galileo, but which were not widely recognized. As Mark Lance writes, “A proposed rule has normative authority because we can show it—within a practice governed by norms of rational argumentation—to be a good one to impose. It has normative authority over us—and already had it over our past uses— . . . because we already had allegiance to a practice within which, as we now know by having done it, it is possible to defend such a rule.”¹⁸ Thus, reasons are in an important sense *timeless*; their truth and their normative binding-ness does not depend on the time of their utterance. If at time t , it is true that “ S ought to do A at time t ,” then at time $t + 100$ years it is true that “ S ought to have done A at time t .” Indeed, we still have a decisive reason to do A even if no one *ever* thinks of R . B would have been a bad course of action even if no one ever thought of R . Reasons, then, can be decisive even if they remain un-uttered. I will discuss in a moment how this conclusion regarding the timelessness of reasons is consistent with my above comments regarding the temporal relativity of the rules of basketball as displayed by the earlier goal-tending example.

As I argued above, the normativity of moral and epistemic discourse is grounded in social practice. But moral truth has some autonomy from social practice in that decisive reasons remain decisive even when they go un-uttered or un-thought. People who don’t think of the best reason, and who go on to perform an inferior action, are mistaken, and the action they pursue is inferior to the one recommended by the unuttered decisive reason. This “timelessness of reasons” will allow us to conclude, in a moment, that moral truths are weakly socially dependent. Their truth is not merely a matter of what most people take their truth to be, although morality is a socially instituted practice. In this instance, *socially dependent* does not entail *relative* or *historicist*. For as we have seen, reasons have an autonomy from our practice of reason-giving, and it is this autonomy that makes the social dependence of reasons weak rather than strong.

We now have two important results. First, morality is grounded in our social practice. Second, moral claims are licensed by the giving of reasons, a fact that gives moral discourse some autonomy from our actual social practice: a reason can be decisive before it is uttered, or hours, days, or centuries after the decision is made, or even if never uttered (or even thought of) at all. When we combine these two conclusions, we get an interesting result. Although moral truth is dependent on social practice, it need not be conceived of as local. Even in societies where slavery is accepted, slavery is still wrong; there are still decisive reasons to abolish slavery (even if no one, not even the slaves themselves, think of these reasons). There might now be a decisive reason for or against abortion; there might be a decisive reason why we should or should not allow gay marriages; or why we should all be vegetarians. These reasons might exist, even though we don’t know it. Thus, moral truth, while possessing some important ties to social practice, need not be thought of as relative

or local. This is because reasons mediate between moral truth and social practice, and keep the connection between the two weak enough to avoid the relativism found in areas of practice (such as word pronunciation) that are more strongly socially dependent.

We must return to the disanalogy between morality and games such as basketball. Recall that basketball is relative in a particular way: before the introduction of the Wilt Chamberlin rules, goal-tending was permissible, and afterwards it was a foul. We should be reluctant to regard morality as relative in the same way. But basketball is relative in this way because our reasons need to be reified to some extent to make the game of basketball work. Considerations of fairness recommend that the rules be standardized for all games, and the best way to do this is to reify the rules.

I should note that even games such as basketball don't absolutely need to be relativistic in the way described. Imagine a group of children playing basketball on the playground. Player *A* has the ball, and says to player *B* on the opposing team, "Uh-oh; here comes your mom, and she looks really mad," when in fact player *B*'s mother is nowhere in the area. While *B* is distracted, *A* drives to the basket and scores. We can well imagine *B* saying, "That wasn't fair;" and we can further imagine our group of playground particularists agreeing that it wasn't fair, and demanding a "do-over." Player *A* could argue that there's no rule against such deceit; but considerations of fair play compel the other players to decide against him. Or we can imagine the players saying at some later date, "We shouldn't have allowed that play to stand." Thus, games such as basketball needn't be reified as they are.

The important lesson to be gleaned from this is that the decision whether or not to reify the rules of a game, and therefore make it historically relativistic, is itself a decision made on the basis of reasons. In the National Basketball Association, virtually all would agree that fairness demands that all play by the same set of rules, and this is best assured by reifying the rules, and modifying them as reason demands. On the playground, though, the same considerations of fairness could allow for punishment or re-play even when no explicitly reified rule forbids the action in question. Similarly, we could argue that we ought not play the moral game relativistically, because we want people to do what they have the *best* moral reason to do; we don't want to absolve of responsibility people living in evil societies (such as Nazi Germany, or the antebellum American South) for the actions they perform, such as genocide and slave-holding.¹⁹ We don't want people to be able to use as an excuse, "Everyone around me was doing it, too." We do not want people to be forced (morally) to do what everyone else is doing, because what everyone else is doing is often bad. Another good reason not to reify moral reasons is that few would submit to a single reifying authority, a problem that does not arise in the formation of new sports leagues.²⁰

Notice that once Wilt Chamberlin became adept at goal-tending, but before goal-tending was disallowed, the fact that goal-tending was too easy for players like Chamberlin was still a good reason to prohibit goal-tending. However, because the game being played was a reified one, this reason was not in itself a sufficient reason for referees to penalize the practice. The rule had to be explicitly reified first. However, in a non-reified game, the fact that a decisive reason obtains for punishing an activity is (or can be) a sufficient reason for punishing that activity.

We thus have some important conclusions. First, moral claims are licensed by the giving of reasons. Second, reasons display a sort of timelessness. Third, morality is a non-reified game. These three conclusions together entail a fourth: it may sometimes be permissible to punish or condemn activities that have not been explicitly prohibited by current moral practice, or that are even condoned by current moral practice. So the mere fact that a game is rule- or reason-constituted doesn't entail that the game must be relativistic.

As I said at the beginning of this section, I spoke only of moral reasons for ease of phrasing, but the conclusions reached in this section apply equally to other sorts of reasons as well: in all cases—moral, epistemic, semantic, and so forth—reasons provide a buffer between social practice and the truth of normative judgments. Thus, a social practice account of normative judgment can avoid attributivism and even historical localism.

IV. DO MORAL CLAIMS NEED TRUTH CONDITIONS?

The above account of the timelessness of reasons also lends itself to another, incompatible account of normative discourse. This account, suggested by Mark Lance and John O'Leary-Hawthorne,²¹ rejects the need for truth conditions for normative discourse, and instead focuses on assertibility conditions. In other words, they reject the need for a metaphysics of normative utterances (what they call a "metaphysically robust truth-conditional theory of normatives") in favor of an epistemology of such claims: if we can get clear on how we successfully justify such claims, we can remain quietists or deflationists about truth. Let us examine this account, and the reasons why Lance and O'Leary-Hawthorne feel driven to it, in more detail.²²

Throughout this paper, I have been giving a social practice account of normative discourse. Normative facts are not objects in Plato's Third Realm; nor are they natural facts whose discovery compelled us to begin practicing morality. Rather, morality and other evaluative practices are a creation of agents; they exist as part of our social practice. But if this is the account of normativity we endorse, then (argue Lance and O'Leary-Hawthorne) any attempt to give the truth conditions for normative claims will lead us headlong into relativism.²³ Lance and O'Leary-Hawthorne write:

By giving truth conditions for normatives as part of a comprehensive theory of the truth conditions of the language and, further, doing so in a way which couches these truth conditions in the terms which are basic in one's ontological conception, one is showing how to reduce the content of normatives to one's favorite ontologically prior notions. . . . Now, if one holds such a reductionist position concerning normatives and is further committed to the claim that the relevant, metaphysically respectable features of the world are social practices, it is hard to see what truth conditions she could give for a claim that one ought to do x other than that one ought to do x if and only if the standards of the relevant social practice license x (or would license x in the long run, or would license x if they deliberated long and hard and rationally, or . . .). To claim this, then, is to deny the possibility that the community's practices might be wrong (or at least to deny that what the community would say in the long run, or would say if they deliberated long and hard . . . might be wrong).²⁴

The problem, then, is that social practice accounts of normative utterances seem to lead straight to relativism. Kripke diagnoses this problem as arising "specifically from the conjunction of the idea of practices as at the bottom of the normative and the idea that normative claims have truth conditions."²⁵ Commitment to the former idea seems to entail that the truth conditions at issue in the latter idea must be in terms of the community's social practices; therefore, "I ought to do X " is true if and only if the relevant social practice licenses X , or would in the long run, etc. According to Kripke, Wittgenstein escapes relativism by rejecting the second of the two conjoined ideas. In other words, he denies that normative claims have truth conditions. Instead, there are justification conditions for normative assertions. Kripke writes, "Wittgenstein proposes a picture of [normative claims] based, not on *truth conditions*, but on *assertibility conditions* or *justification conditions*."²⁶ Thus, we need not be committed to thinking that X is right merely because the underlying practice endorses X . If I can successfully justify the claim "We ought not do X ," then the community is obligated to revise the underlying practice.

It may be thought, though, that this solution only pushes the relativism worries back a step. After all, since we have concluded that all normative utterances must be interpreted in terms of social practices, we must give a social practice account of the community's epistemic norms, as well. Thus, "[w]e will have escaped a metaphysical tie between the community's practices and the *truth* of normative claims, only to be confronted with an equally strong tie between the community's practices and the *justifiability* of such claims. The difference is not one to write home about."²⁷ Let us consider an example to illustrate this point. Let us suppose that a community has a practice of slaveholding. On a naive social practice account, we might think that the claim "Slavery is permissible" is true if and only if the

community's practice permits slaveholding. However, following Kripke, we have rejected this naive social practice account in favor of one that looks for justification conditions, not truth conditions. This leads us to examine the epistemic practices of the community. Suppose the community in question holds the Bible to be an epistemic authority; it is consulted in matters as diverse as morality and science. This highlights the problem with which we are faced: we attempted to escape relativism by moving from truth conditions to justification conditions, but it seems as though justification will itself be a relative matter; *X* is justified if and only if it is supported more strongly than any other belief or course of action by the community's epistemic standards. If the community's epistemic practice is structured according to careful reading of the Bible, then that seems to be the only available standard of justification. And since the Bible seems clearly to endorse the practice of slavery, then (in our hypothetical community) the claim "Slavery is permissible" seems justified.

There is an escape from this sort of relativism, as well. We have already seen that the community has in place mechanisms for revising its practice. If you can justify the claim "We ought to do *X*," then the community is committed to revising its practice so that it includes doing *X*, if the practice does not already endorse the performance of *X*. But there is no reason in principle why the community could not also have in place mechanisms for the revision of the epistemic standards in question. One could, for example, point out contradictions in the Bible,²⁸ and use these to undermine the epistemic authority of the Bible. Or one could challenge the moral authority of the Bible by pointing out some of its odder moral injunctions, like the prohibitions against cattle cross-breeding and the wearing of mixed-fiber clothing.²⁹ Thus, we need not accept any direct entailment from what the practice regards as justified to what is, in fact, justified. These epistemic standards can be just as revisable as the first-order normative judgments (e.g., "We ought to do *X*") themselves.

Indeed, we need regard nothing as in-principle unrevisable. Some claims may be *de facto* unrevisable. For example, it is difficult to imagine what could possibly license revision of the rules of arithmetic. The important lesson, though, is that nothing is *de jure* unrevisable—not a community's moral claims, not its epistemic standards, not even its standards for revising these epistemic standards. As Quine points out in "Two Dogmas of Empiricism,"

Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the fact of recalcitrant experience by pleading a hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying

quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?³⁰

Thus, nothing need be treated as *de jure* unrevisable. Indeed, there is good reason to think that we *ought* to treat nothing as *de jure* unrevisable. As Lance and O’Leary-Hawthorne write, “A practice could in effect adopt the positivist proposal of treating a whole bunch of claims as *de jure* unchallengeable (if the challenge is to their truth). But we would not be tempted to adopt such a practice ourselves. Such a practice seems to encourage—even be constitutive of—dogmatism, preclude dialogue, induce cognitive sterility, and all at no obvious gain.”³¹ We can, therefore, give a social practice account of norm-governed behavior that does not lapse into crude cultural relativism.

Practices, then, are the foundation of justification. They are not foundational in the sense of certain or unrevisable, though. Instead, these practices are the first level of a hierarchy. The second level consists of explicit normative judgments about the appropriateness of the practice (e.g., “One ought not hold slaves”). These utterances can serve to endorse the practice, or to call for its revision.³² Of course, such utterances must be adequately justified if they are to warrant revision of the underlying practice; the mere ability to assert a normative claim does not entail entitlement to that claim. This brings us to the third level: the epistemic standards (say, the authority a practice accords to appeals to sacred texts, or to prophets, or to double-blind studies, or to Tarot cards) that determine whether a second-level claim has or has not been adequately justified. Even these standards can be challenged, though; at the fourth level are standards for the revision of the given epistemic standards. For example, as I noted above, one might challenge the authority accorded the Bible by pointing out contradictions; similarly, one might challenge the authority of Tarot cards by conducting a double-blind study which shows that the accuracy of Tarot card readers is at the level of pure chance. Even these standards can be challenged. Indeed, there is no level at which one must endorse the claim, “My practice endorses X; therefore, X is right.”

To head off any misunderstanding, let me make the following clarification: we have seen that the Lance and O’Leary-Hawthorne picture of the normative has the following “levels”:

1. Social practice.
2. Explicit normative claims endorsing or revising practice.
3. Epistemic standards.
4. Standards for revision of epistemic standards.

The required clarification is that the relation among these levels is not linear; rather, it is one of interdependence. If it were linear, then a claim at one

level could *only* be revised by a claim at the next level. For example, I could only challenge a level 2 utterance (“Slavery is wrong”) with a level 3 claim showing that I did not use the appropriate justificatory method to justify my level 2 claim. This seems to lead to an infinite regress of levels and rules. But this is not how our game of giving and asking for reasons proceeds, or ought to proceed. After all, I can challenge a level 2 claim with another level 2 claim. For example, I might try to stop you from releasing my slaves, arguing that I signed a contract to buy them and paid money for them, and you can’t just liberate against my will something I paid for fair and square. Here, I am appealing to standards of ownership and private property. You can challenge this with numerous other level 2 claims: cruelty is wrong; people have the right to self-determination; and so forth. In this case, the moral claims you advance defeat the moral claims I advance. Thus, we have a case where a level 2 claim is used to defeat another level 2 claim. To use another example, I might use a level 2 claim to overthrow a level 3 claim. Consider a level 3 claim such as, “The Bible is the ultimate authority on matters of religion.” I might use a level 2 claim—“Slavery is morally abhorrent”—to argue that, since the Bible endorses slavery, the Bible cannot be regarded as a moral authority. The same thing happens in science: I can use a level 2 claim (e.g., light travels at 186,000 miles per second) to challenge an epistemic method, which is a level 3 standard. If this standard yields answers incompatible with what we know about the speed of light, then we have reason to question the authority of this standard. Thus, Lance and O’Leary-Hawthorne need not be committed to an infinite regress of rules.

Another thing to notice is that social practice doesn’t merely underlie our level 2 claims; the social practice underlies all levels. Any question of how to go on as we have before, or how correctly to apply a standard (at any level) will be answered in terms of social practice. So the picture emerging of our practice and its standards for revision is looking less like a rigid hierarchy and more like a web of interlocking beliefs, standards, and methods, with the social practice conferring meaning on all of these.

This picture may seem to confine all standards of correctness to within our practice. I hope it seems that way, because that is how things actually are. Questions of belief and rationality are questions to be answered “according to the standards we have at our disposal (barring emendation of them).”³³ *Of course* our standards will be our standards; whose could they be, if not ours? And *of course* these standards will be internal to our practice; it is not coherent to speak of judging in terms of standards that are not ours. As Putnam has put the point, “Well, we should use someone else’s conceptual system?”³⁴ Of course, we can revise the practice, if we think it would be rational to do so, but the question of whether and how to revise it are questions that arise within the practice, as well. You must take the

practice as your starting point. If this be relativism, then relativism is at the center of all standards of rationality, moral and scientific, for all questions of what it is rational to believe arise within our practice.

We've seen above how this social practice account need not commit us to attributivism: since any standard in the system can be challenged, we need never concede that a given normative standard is true merely because it is endorsed by the community. However, in giving up on truth conditions, one might think we are giving up the truth-aptness of normative discourse. Fortunately, abandoning robust truth conditions for normative utterances does not mean we cannot claim, of individual normative utterances, that they are true or false. Indeed, the sort of view endorsed here lends itself nicely to a deflationary account of truth. According to such an account, the sentence "*p* is true" has no content over and above what is had by the sentence "*p*." This is the approach taken by Lance and O'Leary-Hawthorne. They attempt to provide not a metaphysics of normatives, but instead an epistemology of normatives, and to this epistemology of normatives, they couple a deflationist account of truth.³⁵

What *makes it true* that it is wrong to treat racial minorities with less dignity than those in the majority? Why, the fact that it is wrong to treat racial minorities with less dignity than those in the majority, of course. How do we know that this is a fact? If a reader would really like to dispute our entitlement to this claim, we could digress into a discussion of the politics of racism, the speciousness of arguments for morally significant biological distinctions along racial lines, etc., giving many arguments which are quite well known. But if you have nothing to say to these, *then you grant* that we are entitled to make the claim. *From this it follows* that we are entitled to claim that it is true. What makes it true that a metal conducts electricity is the presence of free electrons within the atomic structure of the metal. What makes it true that racism is wrong, is the equal dignity deserved by people irrespective of race. We simply deny the need for any deeper metaphysical grounding.³⁶

V. WHICH ACCOUNT OF MORAL TRUTH SHOULD WE PREFER?

The reader might be left wondering whether we should prefer (a) Wiggins and McDowell's account of moral truth in terms of excellence of reasons, or (b) Lance and O'Leary-Hawthorne's denial that normative claims have truth conditions, being more akin to legislation than description. As we have seen, Wiggins and McDowell's account can easily be incorporated within a social practice framework; Lance and O'Leary-Hawthorne are advancing an explicitly social practice account of normative utterances. Further, the account I have given of the timelessness of reasons can serve

Lance and O’Leary-Hawthorne’s account of moral reason-giving just as well as at can serve Wiggins and McDowell’s account of moral truth.

Guidance is to be found in Lance and O’Leary-Hawthorne’s warning that any attempt to give what they call a “metaphysically robust truth-conditional theory of normatives” leads directly to relativism. Ultimately, this consideration favors the account given by Lance and O’Leary-Hawthorne. This is a point that clearly needs to be argued: although Wiggins and McDowell give truth conditions for moral utterances, it does not seem as though they are giving “metaphysically robust” truth conditions—after all, the truth conditions they give are in terms of excellence of reasons, not in terms of, say, descriptive features of a social practice.

Thus, you might think that Wiggins and McDowell are able to respond to Lance and O’Leary-Hawthorne as follows:

A social practice account of morality need not give the truth conditions of moral claims in terms of what the social practice actually endorses, or would endorse in the future, etc. Our truth-conditional account of moral claims preserves the tie to social practice (in virtue of being in terms of reasons), but claims that the truth conditions are entirely independent of what society believes, or will believe, or would believe under certain conditions. Moral claims are true if the best reasons support them, and there is no reason to suppose that a society currently possesses the best reasons, or will ever possess them all, or even would think of them under certain ideal conditions. The important thing to keep in mind about our account is this: *our current moral theory is precisely only a theory. It is a theory of what the best reasons would support.* We don’t know if the reasons we now employ are the best reasons, or if there are better reasons that elude us. We can test the theory by thinking really hard about moral issues, and by arguing about them, but we need not concede that the mere fact that a particular practice endorses moral conclusion *C* entails that *C* is correct.

Notice that all of the factors that bear on Lance and O’Leary-Hawthorne’s account of the epistemology of normatives also bear on our account of whether reason *R* is the best reason. Lance and O’Leary-Hawthorne claim that their account is not relativistic because all of the standards can be challenged; we are not stuck with the epistemic standards our society *actually* endorses, or the standards for revision our society *actually* endorses. But we can help ourselves to this same notion that everything is defeasible, and avoid relativism as well. In other words, we do not need to endorse moral conclusion *C*, because we can argue that *C* is not supported by the epistemic criteria we share, and is hence not supported by the best reasons. Or we can argue that it is supported by these epistemic considerations (say, the Bible says that *C*), but that these epistemic

standards are themselves flawed. Again, this will allow us to argue that *C* is not supported by the best reasons. So the same considerations that allow Lance and O’Leary-Hawthorne ultimately to escape relativism—the revisability of our norms based on new evidence, the revisability of these very standards of evidence, the revisability of our standards of revision, and so forth—allow our position to avoid relativism as well, because we will never be *forced* to concede the entailment from “Social practice *P* endorses reason *R*” to “Reason *R* is the best reason, and hence the moral truth-making reason.”

In other words, Wiggins and McDowell might respond by saying that a metaphysically robust truth-conditional theory of normatives coupled with a social practice account of moral discourse only leads to relativism if Lance and O’Leary-Hawthorne are right that the truth conditions must be given in terms of some descriptive feature of the underlying social practice. But this assumption is not justified, since Wiggins and McDowell give truth conditions in terms of excellence of reasons instead of in terms of descriptive features of the practice. Hence, one might think that Wiggins and McDowell’s account need not lead to relativism as feared by Lance and O’Leary-Hawthorne.

However, Wiggins and McDowell’s account can only avoid relativism by assimilating itself to the account given by Lance and O’Leary-Hawthorne. Let me explain: Wiggins and McDowell have given the truth conditions for moral utterances in terms of excellence of reasons. But someone who wants the concrete truth conditions for a moral claim will not rest satisfied at that point. Consider sentence *p*, “Cruelty is wrong.” The critic of Wiggins and McDowell will claim, “You have not given me the truth conditions for *p* because you haven’t yet given me the truth conditions for the sentence, ‘*p* is supported by the best reasons.’ I want to know what it *is* for something to be supported by the best reasons.”

At this point, Wiggins and McDowell have two options. They can try to give “metaphysically robust” truth conditions for sentences of the form, “*p* is supported by the best reasons.” But if their account of moral truth is given a social practice rendering, then it seems likely that their account will, in fact, lead to relativism. After all, if the account of moral truth is a social practice account in terms of excellence of reasons, and we want to give metaphysically robust truth conditions for sentences asserting that such-and-such is supported by the best reasons, then these robust truth conditions will most likely be in terms of the underlying social practice. And hence moral truth is reduced to descriptive features of the underlying social practice; the end result of Wiggins and McDowell’s theory is attributivism.

On the other hand, Wiggins and McDowell might deny that sentences of the form “*p* is supported by the best reasons” have metaphysically robust

truth conditions. But then their account is assimilated to Lance and O’Leary-Hawthorne’s account, for they have claimed that at bottom no robust truth conditions can be given for normative utterances. If Wiggins and McDowell are going to say that “ p is supported by the best reasons” is true if and only if p is supported by the best reasons, then it is not clear what remains to distinguish their account from the one advanced by Lance and O’Leary-Hawthorne.

VI. CONCLUSION

Some skepticism about social practice accounts of normativity no doubt stems from what Rorty calls “the religious need to have human projects underwritten by a nonhuman authority.”³⁷ This way of thinking underlies our intuition that science is the paradigm of objectivity; science is objective because it is made true by something non-human. I suggest that this intuition should be construed as follows: robust objectivity requires that the truth of a type of discourse be independent of what we take it to be. This is why we think science is the paradigm of objectivity: there is a world out there that (in some sense) makes scientific truth independent of what we take it to be. But morality and epistemology have this same virtue: moral and epistemic truth is also independent of what we take it to be. In this respect, then, morality and epistemology are as objective as science. And I suspect this is the most interesting kind of objectivity, in any case.

ENDNOTES

I am extremely grateful to Mark Lance, Maggie Little, and Henry Richardson, who read numerous earlier drafts of the original paper and provided the crucial feedback that gave the paper its current form. Mark Murphy and Steve Kuhn also read an early draft, and provided excellent constructive criticism.

1. I owe the terms “transcendentalist” and “attributivist” to Lance, Mark, and John O’Leary-Hawthorne, *The Grammar of Meaning: An Exploration of the Normative Character of Semantic Discourse* (New York: Cambridge University Press, 1997).
2. *Ibid.*, 191.
3. *Ibid.*
4. *Ibid.*, 193.
5. *Ibid.*
6. There are at least two existing moral theories that are neither attributivist nor transcendentalist: emotivism and prescriptivism. But neither of these can make sense of the way in which we provide evidence for our moral claims. The emotivist claims that moral utterances are mere expressions of emotion, and hence cannot account for the way in which we provide evidence for moral claims. After all, nothing can stand in an

evidential or inferential relation to an interjection. The prescriptivist, according to whom moral utterances are commands, also has a problem with evidence: when defending moral claims, we justify their content, not their utterance, whereas with commands, we justify their utterance, not their content. For more on these dual errors, see Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, 197–204.

7. *Ibid.*, 176.

8. *Ibid.*, 175–176.

9. David Wiggins, "Moral Cognitivism, Moral Relativism, and Motivating Moral Beliefs," *Proceedings of the Aristotelian Society* 81: 66–85.

10. John McDowell, "Critical Notice of *Ethics and the Limits of Philosophy*, by Bernard Williams," *Mind* 95: 377–386.

11. Wiggins, "Moral Cognitivism, Moral Relativism, and Motivating Moral Beliefs," 67.

12. *Ibid.*

13. Gilbert Harman, "Ethics and Observation," in *Essays on Moral Realism*, ed. Geoffrey Sayre-McCord (Ithaca, N.Y.: Cornell University Press, 1988), 119–124.

14. Wiggins, "Moral Cognitivism, Moral Relativism, and Motivating Moral Beliefs," 66.

15. On the pragmatist line I advocate, pragmatic considerations justify a practice, and individual actions are justified by appeal to the rules constitutive of the practice. The locus classicus for this view is John Rawls, "Two Concepts of Rules," *The Philosophical Review*, 64.1 (1955): 3–32. I explore the pragmatic justification of practices in Jeremy Randel Koons, "Do Normative Facts Need to Explain?" *Pacific Philosophical Quarterly* 81: 246–272.

16. Some people have disputed my claim that pronunciation is strongly socially dependent in this way. If they are right, then my case is hardly weakened: the less plausible it is to claim that pronunciation is strongly socially dependent, the less plausible it is to claim that morality is strongly socially dependent.

17. I argue for this pragmatist conception of morality in Koons, "Do Normative Facts Need to Explain?"

18. Mark Lance, "The Word Made Flesh: Toward a neo-Sellarsian View of Concepts, Analysis, and Understanding," *Acta Analytica* 15: 117–135, 128.

19. This is perhaps a bit carelessly stated. As noted earlier, we should distinguish between whether an action is objectively right, and whether the person was (given her historical context) justified in thinking the action was right. Thus, even if, say, premarital sex is permissible, a Northern European in the year 1600 A.D. might well have been justified in thinking that it was immoral.

20. Although some communities no doubt regard moral rules as reified in this way. Certain elements of Catholic practice, for example, resemble reified morality. Eating meat on Friday used to be permissible, but became immoral. Similarly, a divine command theorist might regard the rules of morality as reified. Indeed, a Christian who was a biblical literalist might, in viewing the vastly different moralities expressed in the Old and New Testaments, conclude that warlike behavior used to be permissible, whereas now only pacifism is permitted.

21. Lance and O'Leary-Hawthorne, *The Grammar of Meaning*.
22. I will present only a brief summary of Lance and O'Leary-Hawthorne's pragmatist account of normative judgment. For the whole story, see *The Grammar of Meaning*, chap. 3.
23. I am grateful to Mark Lance for raising this objection in conversation, and for directing me to the relevant portions of his book.
24. Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, 210.
25. Mark Lance, "Rules, Practices and Norms," in *Ludwig Wittgenstein: A Symposium on the Centennial of His Birth*, ed. Souren Tegrarian, Anthony Serafini, and Edward M. Cook (Wakefield, N.H.: Longwood Academic, 1990), 77–86, 82.
26. Saul Kripke, *Wittgenstein on Rules and Private Language* (Oxford: Basil Blackwell, 1982), 74.
27. Lance, "Rules, Practices and Norms," 83.
28. Compare, for example, the following passages: Genesis 6:19–20 vs. Genesis 7:2–3; 1 Samuel 17 vs. 2 Samuel 21:19; 2 Chronicles 36:9 vs. 2 Kings 24:8; 2 Samuel 24:1 vs. 1 Chronicles 21:1; Mark 11:7 vs. Matthew 21:7; also, the different order of creation in Genesis 1 and 2.
29. See Leviticus 19:19.
30. W. V. O. Quine, "Two Dogmas of Empiricism," in Quine, *From a Logical Point of View* (New York: Harper Torchbooks, 1953), 43.
31. Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, 120.
32. This interpretation of second-level utterances is one of the novel features of the account given by Lance and O'Leary-Hawthorne (see Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, especially chapters 1 and 3). Most thinkers regard the purpose of such utterances as descriptive. Brandom, for example, thinks that such utterances serve to codify or make explicit features of the underlying social practice. See Robert Brandom, *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Cambridge, Mass.: Harvard University Press, 1994).
33. Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, 237.
34. Quoted in John McDowell, "Aesthetic Value, Objectivity, and the Fabric of the World," in *Pleasure, Preference and Value*, ed. E. Schaper (Cambridge: Cambridge University Press, 1983), 1–16, 14.
35. To be precise, Lance and O'Leary-Hawthorne endorse the theory of truth presented in D. L. Grover, J. Camp, and N. D. Belnap, Jr., "A Prosentential Theory of Truth," *Philosophical Studies* 27: 73–125.
36. Lance and O'Leary-Hawthorne, *The Grammar of Meaning*, 226.
37. Richard Rorty, *Contingency, Irony, and Solidarity* (New York: Cambridge University Press, 1989), 52.