

Ethical Internalism and Glaucou's Question

SERGIO TENENBAUM
University of New Mexico

Ethical internalism, roughly, affirms a certain specific conceptual connection between taking X to be morally right (or X being morally right) and accepting that there is at least a *prima-facie* reason to do X (or there being a *prima-facie*¹ reason to do X).² Ethical externalism, as I will understand it, denies the existence of a conceptual connection between taking X to be morally right (or X being morally right) and accepting that there is a reason to do X (or there being a reason to do X).³ In this paper, I first briefly clarify the versions of ethical internalism and externalism with which I will be concerned. Next I examine an important criticism that has been raised against internalism. Internalism seems to be unable to make much sense of the possibility of a question which I will call 'Glaucou's question':

"I see that X is morally right, but is there is a reason to do X ?"

It seems that an internalist must be committed to the view that Glaucou's question is empty.⁴ I then briefly examine further difficulties that internalism seems to face.

These problems seem to cast serious doubt on the plausibility of internalism. However, I argue that ethical externalism is also unsatisfactory. I proceed then to define and defend a view, which I call 'weak internalism', that does not fit either of these categories. In fact, if this view is correct, the dispute between internalism and externalism has rested on an overly narrow view of the possible relations between a concept (such as 'morally right') and an assumption (such as "if X is morally right, then there is a reason to do X "). I suggest that the alternatives have so far been mapped out incorrectly as a consequence of implausible views about the nature of ethical concepts. Once we correct these distortions, we can see that

'weak internalism' can accommodate the plausible insights of externalism without rejecting the conceptual connection between something being morally right and its providing a reason to act.

I

We can distinguish between at least two kinds of conceptual connections one can find between moral predication and reasons to act. One might think that an *agent* who judges that *X* is morally right must thereby also judge that she has a reason to do *X*. Or one might think that the *fact* that *X* is morally right gives a reason for the agent to do *X*, whether or not the agent recognizes such a reason. Following Stephen Darwall,⁵ we can call the first kind of internalism, "judgment internalism", and the second kind, "existence internalism". These internalists are committed, respectively, to the following claims being *conceptually* true:

Reason-Giving_j Assumption: If one accepts that *X* is morally right, one thereby accepts that there is a reason to do *X*.

Reason-Giving_e Assumption: If *X* is morally right, then there is a reason to do *X*.

When I use in this paper the unsubscripted expression "reason-giving assumption" I mean to refer to the disjunction of these two subscripted reason-giving assumptions. Although different internalists may understand differently what is implied by saying that this relation is a matter of a conceptual connection, they would all agree that the relation is necessary; indeed, an externalist will typically accept that it is contingently true that we (or at least most of us) have reason to pursue what is morally right.

There are undeniably some advantages to the judgment form of internalism. First it does not commit us to decide between claiming that nothing is morally right and rejecting the rather popular instrumental theory of rationality. For if existence internalism is true, then if *X* is morally right, then there is a reason to do *X*. But if we assume that this reason is unqualified and thus holds for all agents, this would be a reason that all agents have independent of the ends that they happen to have, and thus an instrumental theory of rationality would be refuted.⁶ It is consistent with judgment internalism that one accepts that *X* is morally right only when one already has the relevant end that gives one a reason to do *X*. Existence internalism does not, strictly speaking, imply judgment internalism, but if the former is true, and especially if it is supposed to be a conceptual truth, an agent who accepts that *X* is morally right, but does not accept that *X* gives her a reason to act seems to suffer from at least an imperfect grasp of the concept. So it is natural to regard judgment internalism as the weaker claim, and thus when discussing internalism, I will focus on this weaker claim. However, the weak internalism I advocate does not require that either reason-giving assumption is

conceptually true,⁷ but it makes use of the existence version of the reason-giving assumption. The reasons for making use of this version as well as the similarities and dissimilarities between weak internalism and the more traditional forms of internalism will be discussed later.

Since the debate between internalism and externalism as presented here is a debate about whether certain conceptual relations obtain or not, it is worth looking into what is implied by this notion. Fortunately my purposes do not require a full account of what a conceptual connection is. At the very least, to say that, for instance, the reason-giving_j assumption is conceptually true is to say that any use of the concept “morally right” that violates this assumption is unintelligible. So it would be unintelligible to say that Mary accepts that charity is morally right but does not accept that there is any reason for engaging in any act of charity. And if one can show that this is an intelligible claim, one has thereby shown that the reason-giving assumption is not conceptually true. Although it is often not easy to say whether a certain claim is intelligible or not, I hope the cases discussed here will be clear enough.

One could, no doubt, define (judgment) internalism and (judgment)⁸ externalism in such a way as to make them contradictory theses, and thus leaving no middle ground between these views. We could indeed define ethical externalism just as the negation of ethical internalism:

It is not the case that if one accepts that *X* is morally right, one thereby accepts that there is a reason to do *X*.

In this case, my view would fall squarely on the externalist side of the debate since the above claim is one that I will end up endorsing. However, I will define externalism as denying *any* conceptual connection between accepting that something is morally right and there being reason to do *X*. This is motivated by considerations of the recent relevance of these issues to the question of moral realism. A traditional argument against moral realism argues that moral realism is incompatible with an instrumental conception of rationality.⁹ According to this argument, generally attributed to Hume, since moral realism requires that certain states of affairs give reasons to agents that are categorical—that is, that are independent of their desires—moral realism conflicts with an instrumental conception of rationality.¹⁰ A recent line of defense of moral realism has relied on externalism to undermine this argument against moral realism. If it can be shown that there is at most a contingent connection between moral properties and reasons to act, one can uphold the instrumental conception of rationality, and still accept moral realism.¹¹

But if externalism has this task to perform, it is important that the connection between morality and reasons for action is wholly contingent; the validity of the moral system should leave open the possibility that we have no reason to act morally. Moreover, according to this picture, the extensions of moral concepts are determined by some natural facts,¹² and the question of whether we should be

motivated by those things that fall under the extension of moral concepts is to be assessed later by examining the desires and pro-attitudes that an agent happens to have. Thus, this form of moral realism is committed to the possibility that one can determine the extension of moral concepts independently of considerations pertaining to the role of these concepts in providing motivations or reasons for actions. That is, the question of whether *X* is morally right, for instance, should not be settled (even in part) by asking first, whether we could be motivated to do *X* or by asking whether I could have a reason to do *X*. Rather questions of motivation and reasons for action enter on the scene only after we have determined what is morally right and we are interested in showing that we do have reasons (or we should be motivated) to do what is morally right. It is in view of this recent history of externalism that I decided to define it in a more stringent manner, as denying any kind of conceptual connection between the reason-giving assumption and moral concepts. And here again, in order to avoid the murky issue of the nature of conceptual connections, it is worth noting that, for the purposes of this paper, all that is required is that externalism be committed to *at least one* of the following claims.

- (a) All non-trivial relations between *X* being (or being judged) to be morally right and the existence (or the judgment that there is) a reason to do *X* depend on empirical facts (in particular, facts about our psychology).

That is, if one claims that any interesting relation obtain between the fact that *X* is morally right and the fact that one has a reason to do *X*—such as, for instance, that most people have a *prima-facie* reason to do whatever is morally right—one will necessarily appeal to a certain fact about our psychology—such as, for instance, the fact that typically our sympathetic feelings prevent us from being happy while never pursuing that which is morally right.

- (b) The reason-giving assumption plays no indispensable role in determining the extension of “morally right”.

That is, according to (b), we should be capable at least in principle of finding a method of determining the extension of ‘morally right’ in a way that does not appeal, explicitly or implicitly,¹³ to the reason-giving assumption.

Since the position that I am defending does not underwrite either of these claims, it seems more appropriate to see it as a form of internalism.¹⁴ Indeed, it is part of the aim of this paper to show that this form of moral realism has unduly profited from the fact that Glaucon’s question is intelligible.¹⁵ But irrespective of this aim or any labels, I hope that the debate will profit from clarifying the relation between the reason-giving assumption and moral concepts. However one may wish to classify the view defended in this paper, my claim is that neither internalists nor externalists have hitherto provided a satisfactory account of this relation.

II

The main problem that internalism seems to face can be put as follows: internalism seems to define out of existence a question about the legitimacy of morality, a question that seems to call for a substantive answer.¹⁶ As we said above, if internalism is right, Glaucon's question must be empty. Let us look again at Glaucon's question:

"I see that X is morally right, but is there is a reason to do X ?"

If internalism is correct, when we say that X is morally right, we must have already accepted that there is a reason to do X . Yet, it would seem that this is a question that can significantly be raised. I have called it 'Glaucon's question', because this seems to be the question that Glaucon raises in the *Republic*, and that Socrates tries to answer.¹⁷ It is a question not about whether we know what is morally right or wrong—not about the *content* of morality—but about whether there are *reasons* to do what is morally right—about the *rationality* of morality.

The most obvious answer that the internalist can give in face of this challenge is to try to interpret Glaucon's question as question that, when properly formulated, can be also considered non-empty by the internalist. Hare, for instance, recognizes that Glaucon's question appears to be non-empty, but he thinks the internalist can explain away this appearance by pointing out that we can use 'morally right' in an inverted comma sense.¹⁸ For, sometimes, we can use 'morally right' not to mean 'morally right', but 'that which is generally called "morally right"' or 'that which I feel compelled to do'. Let us define 'Hare-substitute' as follows:

q is a Hare-substitute of an instance i ¹⁹ of Glaucon's question if and only if:

- (i) q expresses the same content as i (roughly, q asks the same question as i).
- (ii) q does not use (although it may *mention*²⁰) a moral term.

The existence of a Hare-substitute for a particular instance of Glaucon's question shows that 'moral' is not being used in its standard sense in this instance of the question. So, for instance, the following might be proposed as Hare-substitutes for different instances of Glaucon's questions:

(Q_1) Why should I do what my friends call 'morally right'?

(Q_2) Why should I do that which, given my upbringing, I feel compelled to do?

We can now say that Hare claims that every non-empty instance of Glaucon's question can be replaced by a Hare-substitute. However, this seems implausible; surely Glaucon's question may be raised seriously by someone who thinks that not all that is generally *called* 'morally right' *is* morally right (since she might think people are confused about some moral issues) and by people who do not

feel any compulsion to do some of the things they call ‘morally right’. Indeed, it is hard to believe that Socrates would find it worthwhile to seek reasons for acting according to what is *generally called* ‘just’. Given his aversion to following the opinion of ‘the many’, Socrates probably thought that Glaucon’s challenge was an important one whether or not most people were right in their opinions about the just.

It is important to note that this problem is not a problem with the Hare-substitutes chosen by Hare himself. The problem in its most general form can be formulated as follows: The internalist claims that there is a conceptual connection between accepting that ‘*X* is morally right’ and accepting that ‘I have a reason to promote *X*’. Once we sever this conceptual connection, we are no longer talking about a proper application of the concept. But it seems that Glaucon’s concern is a concern about ‘morally right’ in its proper application; that is, that we might have no reason to be moral *in the proper sense of ‘moral’*. Any reinterpretation of the word ‘moral’ will not be addressing the same concern as the one expressed in Glaucon’s question.

There seem to be further considerations that lend support to the externalist view. We could think of the externalist and the internalist as defining two different notions of ‘morally right’, one (say, ‘right_i’) that implies that there is a reason to do that which is right_i and another (say, ‘right_e’) that does not have this implication. It seems that ‘right_i’ can be defined in terms of ‘right_e’ as follows:

X is ‘right_i’ if and only if X is ‘right_e’ and there is a reason to do X²¹

However, ‘right_e’ cannot be defined in terms of ‘right_i’. If we try to define ‘right_e’ in the same way (‘X is ‘right_e’ if and only if X is ‘right_i’ but there is no reason to do X’), ‘right_e’ would turn out to be an empty concept since the right side of the biconditional could never be true. So, whatever one’s views about common usage is, it would seem philosophically wise to treat ‘right_e’ as primary, and treat ‘right_i’ as a concept that can be analyzed in terms of ‘right_e’.

Of course the internalist can insist that irrespective of which concept is primary, it is still true that right_i captures the standard usage. But her case would then be much weakened. For after all the internalist does not deny that one often uses the term “morally right” in ways that at least *appear* to violate the reason-giving assumption. So, the availability of this possibility of analysis makes the appeal to Hare-substitutes seem unnecessarily circuitous; we first postulate that ordinary discourse uses a concept analyzable into rather heterogeneous parts as a primitive, and then try to explain away the evidence to the contrary as improper uses of the term. On the other hand, the externalist account will in comparison look rather straightforward: we can say that we always use the term “morally right” univocally, and the externalist may even concede that we often assume that we care about morality, and thus often take it for granted the truth (albeit not the conceptual truth) of the reasons_j-giving assumption.

If the internalist wants to make his position plausible, he must explain what goes wrong if we take the concept ‘right_e’ to be basic in this manner. The most

straightforward way of doing this is to show that the concept ‘right_e’ is not available as a primitive concept. But if he also wants to preserve the intuition discussed above he must at the same time show why this unavailability of the concept ‘right_e’ does not make Glaucon’s question empty. This is indeed what I propose to do in the following sections, but before we move on, we should first examine a proposal that promises to make room for Glaucon’s question without replacing it with a Hare-substitute.

Michael Smith proposes another way to understand Glaucon’s question within the internalist framework. Although his strategy would also take the non-empty form of Glaucon’s question to be using an inverted comma sense of ‘morally right’, it does not make use of any particular Hare-substitute.²² He compares the person who asks Glaucon’s question with the blind person who seems to have mastered color concepts. Some philosophers argue that employing a color concept requires having had the appropriate experience.²³ Such theorists are committed to assigning a different meaning to the blind person’s utterances of words such as ‘blue’ and ‘green’, even if the blind person turns out to be as competent as any sighted person (or even, thanks to some spectacular prosthetic device, *more* competent) in determining whether objects are blue or green. Smith points out that one cannot refute such philosophers just by saying that they are not taking seriously enough the blind person’s utterance of “grass is green”; it is not immediately obvious that this can be considered a counter-example to the theory without further articulation. Similarly, Smith suggests that the person who rejects the reason-giving assumption, whom he, following Brink, calls the ‘amoralist’, might be using a concept that is extensionally equivalent to our concept ‘morally right’, and thus might be even better than ourselves in determining the extension of the ordinary concept ‘morally right’. But the amoralist will not necessarily be using a concept that has the same *content* as the ordinary ‘morally right’. So one cannot, without further argument, use this possibility as grounds to reject the claim that the reason-giving assumption is part of the content of (the ordinary concept) ‘morally right’.

Although Smith does not draw attention to this point, the plausibility of this defense of internalism depends on our learning from the analogy that the internalist does *not* need to provide a Hare-substitute. In the color case, we might not be able to find in our vocabulary any equivalent to the concept that this blind person expresses when she says ‘blue’.²⁴ It certainly cannot be rendered as ‘what most people would call “blue”’, if we assume that this blind person can become even more competent at discriminating blue objects than most sighted people are. But even if this were the case, we are not entitled to conclude from this fact that her utterances of ‘blue’ have the same meaning as the utterances of someone who has visual experiences, or that the possession of our concept ‘blue’ does not require the appropriate visual experience.

Once we see this point, it seems that Smith’s suggestion provides a better account of Glaucon’s question than any account that attempts to replace all serious instances of the question with Hare-substitutes. We can see this point better if

we look at Brink's complaint against Hare's account of the person who asks Glaucon's question:

Why should we assume that the person who asks 'Why should I be moral?' is using moral language in inverted commas or is mistaken about what morality requires? Why can't someone have *correctly identified* his moral obligations and still wonder whether these obligations give him good reason for action?²⁵

But according to Smith the internalist can accept that the person who asks Glaucon's question *might have correctly identified her obligations*, or, in our language, she might have correctly identified the actions that are morally right. Even if she rejects the reason-giving assumption and thus could not be using a term that has the same content as the ordinary 'morally right', she might still be using a concept that is extensionally equivalent to the ordinary 'morally right'. We may even leave open the possibility that the person who asks Glaucon's question is as good as anyone else, or even better, in determining the extension of 'morally right', without accepting that she can master the ordinary concept 'morally right' while rejecting the reason-giving assumption.

Yet, Smith's proposal fails for the same reason that the attempts to replace Glaucon's question with Hare-substitutes fail. Under his proposal, the amoralist who asks Glaucon's question seriously does not use the term 'morally right' in the same way we do. Therefore, Glaucon's question is no challenge to *our* moral behavior, to our capacity to justify doing what is morally right, *in our sense of morally right*. If we understand Glaucon's question in this manner, it becomes all too easy for Socrates to answer it. All he should say is "Perhaps in this sense of 'morally right' we have no reason to do what is morally right; but, in the usual sense of 'morally right' it is conceptually true that we have a reason to be moral."

Moreover, suppose we accept that Smith is right that the content of 'morally right' used by the person who asks Glaucon's question is different from the one we use when we accept the reason-giving assumption. This concept would be much like the concept 'right_e' discussed above, and Smith would have a similar problem trying to explain why we should not treat the concept 'right_e' as primary and use it to define 'right_i'.

Smith might argue that, much as the blind person's concept 'blue' might not be available to the sighted person, there is no guarantee that we, as good moral beings, have available to us the concept 'right_e'. As I suggested above, I will endeavor to show that there is something right about this thought. However, as it stands, it seems that the burden of proof is on the internalist. For the person who raises Glaucon's question does not belong to a subgroup of the population whose perceptual (or any other sort of cognitive) capacities are significantly different from those who pertaining to another subgroup, as in the case of the blind with respect to the sighted person. Much less is the person who raises Glaucon's question necessarily a being whose point of view we, as moral people, could hardly conceive of, someone who is an alien to our conceptual scheme, as this response

would suggest. Glaucon seems to raise this question, Socrates seems to understand it, we raise it to our students in many ethics classes, and we may ask ourselves this question rather seriously when bewitched, even if temporarily, by the charms of self-interested, uncaring behavior. Far from being unavailable to us, the concept seems to show up whenever we reflect upon our ethical life.

This should help us foreclose another strategy that seems open to the internalist. The internalist might think that we are giving too much leeway to the intuition that a certain question should make sense. In light of the overwhelming evidence for internalism (assuming that there is such overwhelming evidence²⁶) we might just accept that Glaucon's question, despite appearances, is confused, and reject the intuition that it is a non-empty question. However, this move would be a plausible one only if we could not reform internalism in such a way that it can accommodate this seemingly important intuition. In what follows I will argue for a reformed version of internalism that *can* accommodate this intuition.

III

It might be worth taking a look at a different kind of concept to which one might take similarly externalist or internalist stances—the concept of Kashrut in the Jewish tradition. Suppose one were to defend an internalist view of 'X is Kosher'. According to this view, accepting that a certain food is 'non-Kosher' entails accepting that there is a reason to refrain from eating it. Here the internalist can use a strategy similar to Hare's to circumvent putative counter-examples. The Kosher-internalist might claim that when people say both that a dish is non-Kosher, and yet there is no reason not to eat it, they must mean something like "this food is generally called 'non-Kosher'". But here the externalist seems to win the day easily. It seems perfectly appropriate to raise a skeptical question about "non-Kosher" foods; in fact most people not only raise the question but reject outright the relation between Kosher judgments and judgments about reasons to eat things. Moreover, it seems that the externalist notion is more basic. For even people who do not accept Jewish dietary laws are capable of classifying foods as Kosher or non-Kosher. So it seems natural to take the externalist Kosher concept, $Kosher_e$, as basic and define the internalist concept, $Kosher_i$, in terms of it as follows:

A certain kind of food is not $Kosher_i$ if and only if it is $non-Kosher_e$ and there is a reason to refrain from eating it.

However, is the externalist notion really primary?²⁷ It is true that in the normal course of events a non-Jew might be as good as an orthodox Jew in classifying food as Kosher or non-Kosher. But suppose now a certain new kind of food appears in the market: a hybrid of a cow and a pig with an artificially modified DNA structure, the 'pig-wow'. How can one tell whether this food is Kosher? I will assume that this genetic wonder is dissimilar enough from a pig that it is no longer obvious that its meat should fall in the same category as pork. We might

imagine that it is part of the agenda of a Rabbinical convention to settle this issue. One could think that this would be a convention on verbal legislation, which will arbitrarily decide how the concept of Kosher should be extended. However, it is unlikely that the convention would simply toss a coin, have some sort of popularity contest, or choose any other arbitrary procedure. Probably the decision would involve close reading of sacred texts, reasoning by analogy from other cases, attempts to figure out a principle suggested by previous decisions and so forth. Let us suppose that after a round of arguments, the line of reasoning that seems to be winning is the following, presented by Rabbi M:

We generally accept that any amount of pig makes a certain food non-Kosher. However, the pig-wow definitely is constituted in part by pig, so it is non-Kosher. Moreover, nothing in any sacred text contradicts this view.

Now we might expect that this argument would not go unchallenged. Some find the analogy somewhat strained. Rabbi L. thinks that if DNA similarity counts as being partly pig, very few foods would count as Kosher, and following this principle would lead to starvation. Rabbi K. does not see why in the absence of any explicit Biblical prohibition we should not be allowed to enjoy the much advertised gastronomic wonders that the pig-wow promises.

It seems that all these forms of reasoning implicitly rely on Jewish tradition, and to determine whether the pig-wow is Kosher, those in the convention must rely on their shared acceptance of the laws and principles embedded in the Jewish tradition. The objections of Rabbi L. and Rabbi K. seem to rely specifically on the inference from the fact that something is non-Kosher to the fact that it should not be eaten. Of course, the externalist may point out that any atheist could have gone through the same reasoning and raised the same objections. Rabbi M. might even be an impostor, an atheist who so loves to participate in these arguments that he pretends to be an orthodox Jew. In fact, he need not be an impostor; well known for his intellectual prowess and knowledge of Jewish tradition, he might have been invited as an independent consultant. For this matter, a practice session of debate contestants might have gone through the same round of arguments, as a mere intellectual exercise. They might be sincerely convinced by the same arguments, and yet take this conclusion to have no dietary consequences.

But if Rabbi M. is to say anything useful in the convention, he has to reason *as if he believed the tradition*, he has to enter the mind of the orthodox Jew, much in the same way that the cynical political strategist can be said to enter the mind of the candid voter. Whether he accepts or rejects Judaism, whether or not he will personally eat pork, when Rabbi M. is reasoning about whether the pig-wow is Kosher, *he must assume that if something is non-Kosher then it should not be eaten*. It would be beside the point if, in response to Rabbi L.'s complaint that his reasoning would lead people to starvation, Rabbi M. were to answer: "I was merely claiming that the pig-wow should be considered non-Kosher. I never meant that we should refrain from eating it." Whether or not he thinks that one should eat only Kosher food, he can't dismiss Rabbi L.'s reasoning in this way.

Let us define a relation between a concept and an assumption as follows (I will call this the relation of ‘a concept being bound up with an assumption’):

A concept is bound up with an assumption if and only if in order to determine the extension of the concept one must suppose that this assumption is true.

Some words of clarification are in order. First, the reason why we need to suppose that the relevant assumption is true is that any valid reasoning that would lead us to rule an object out of the extension of a concept or include it in must be at least compatible with this assumption. Furthermore, we may use this assumption to adduce evidence on whether something should be included in the extension of the relevant concept. So if the concept ‘Kosher’ is bound up with the assumption that one should not eat what is non-Kosher, Rabbi L.’s reasoning is perfectly legitimate and Rabbi M. cannot reject it by saying that he was not suggesting that one should refrain from eating what is non-Kosher.

Secondly, the fact that a concept X is bound up with an assumption p does not mean that one cannot say that something is X while doubting that p is true or thinking that it is false, in the same way as one can suppose that q is true for the sake of argument, without committing oneself to the truth of q . In some cases a person who denies p will thereby not share a point of view from which the concept X might have an important role to play in her life.²⁸ However, even with regard to these cases it is important to note that one might not fully share a point of view, and yet be capable of *understanding* it. And so, in any case, even if we accept that ‘being Kosher’ is bound up with the many assumptions of the Jewish tradition, the person who understands, but does not endorse, those elements of the Jewish tradition could raise the counterpart to Glaucon’s question about Kosher. Similarly, a person who is fully committed to all these elements of the Jewish tradition might, in a reflective mood, stand back from this commitment, and raise the counterpart to Glaucon’s question.

Thirdly, it is important to distinguish between an assumption being bound up with a concept and using an assumption as a heuristic device. If we look back at the Kosher example, the main use of the correlate to the reason-giving assumption was to propose objections to a certain way of determining the extension. One might think thus that it is part of our limited cognitive powers that we need to check purported candidates for being Kosher and non-Kosher against this assumption. Smarter beings might be capable of determining the extension of “Kosher” without any such reliance. Moreover, one might think reflection on this possibility can show that our reliance on the corresponding considerations to determine the extension of ‘morally right’ is dispensable. Although we often need to appeal to those considerations, smarter beings, and even ourselves with the progress of science could determine what falls in the extension of ‘morally right’ without reliance on these considerations. But here we can gain a sense of the difference between a heuristic device for determining the extension of a concept and a consideration that is in fact necessary to

determine the extension of a concept. It might be that, being not very good at deciding whether something is morally right or not, I assume that my uncle Otto always knows which actions are morally right. Being more proficient at these matters, Gretchen doesn't need to appeal to my uncle Otto's judgment. One day Gretchen tells me that *X* is morally right, and I go ask uncle Otto who says that nothing is further from the truth; *X* is certainly not morally right. I go back to Gretchen and tell her that her claim cannot be true because it is incompatible with the assumption that uncle Otto always knows which actions are morally right. Now Gretchen would be certainly in the right if she answered this charge by saying that this cannot settle the issue; uncle Otto might be simply wrong about what is morally right. Using our vocabulary, Gretchen should say that the assumption "Uncle Otto is always right about which actions are morally right" is a heuristic device that I use to determine the extension of "morally right", but it is not bound up with this concept.

Whether or not one explicitly relies on the Kosher version of the reason-giving assumption, no analogous reply is available in this case. Even if certain beings did not need to use the Kosher reason-giving assumption to determine the extension of the concept "Kosher", if this concept is indeed bound up with this assumption, the determination of the extension must be compatible with this assumption. And if, looking at the way they determined the extension of Kosher, we raise the objection that people will starve if they refrain from eating everything that is not Kosher, they would have to regard this as a legitimate challenge, a challenge that would require a response other than "this is true, but we never suggested that one should refrain from eating that which is not Kosher". So even if God could determine the extension of "Kosher" by direct intuition, if it were incompatible with the Kosher reason-giving assumption, and could not answer challenges which relied on this assumption, we would have to regard God's intuition as defective.

Finally, although I have defined 'being bound up' as a relation between a concept and an *assumption*, this relation could be generalized so that a concept could be also bound up with a point of view that includes not only assumptions, but dynamic norms of reasoning, sets of authoritative texts, and so forth. Determining the extension of a concept might depend on assuming that certain propositions are true, that certain forms of reasoning are valid, and that certain texts have some form of special authority. In fact, this is further complicated by the fact that any of these aspects might be challenged. A Jewish sect might defend the view that some supposedly authoritative text is in fact a forgery, and it might end up being capable of justifying its claim to the wider Jewish community. A concept that is part of the Jewish tradition²⁹ is probably bound up not with a single assumption but with the Jewish point of view.³⁰ This is important because it shows that we cannot identify the fact that a concept is bound up with a set of assumptions or a point of view with the fact that a proposition follows from a certain axiomatic system. There might be no way to specify all the assumptions that determine whether something falls in or out of the extension of Kashrut, so that

there might be no non-trivial³¹ axiomatic system that determines whether something is Kosher or not.

In any case, if all the above is correct it seems that neither the externalist nor the internalist got it right about the concept 'Kosher'. The right view is what I will call 'weak internalism':

The concept 'Kosher' is *bound up with* the assumption that if *X* is non-Kosher, then there is a reason to refrain from eating *X*.

It should come as no surprise that I think weak internalism is the right view in ethics. With the appropriate permutations, weak internalism in ethics would read as follows:

The concept 'morally right' is bound up with the assumption that if *X* is morally right, then there is a reason to do *X*.

In other words, weak internalism claims that the concept 'morally right' is bound up with the reason-giving assumption. Note that although the assumption used here is the existence version of the reason-giving assumption, weak internalism is in no sense a stronger thesis than the judgment version of classical internalism. Since the truth of this assumption is not necessarily implied by the correct use of the term 'morally right', but is only bound up with the concept, judgment internalism is not a consequence of it even when we add some other plausible assumptions or assume unbounded conceptual clarity on the part of the speakers.

IV

I will now proceed as follows: I will first explain why weak internalism is significantly different from internalism. I will then go on to explain why I believe weak internalism is a plausible view in ethics. Finally, I will explain why weak internalism is different in important ways from externalism.³²

One might wonder whether weak internalism is not merely a notational variant of full-blown internalism. For, given that the reason-giving assumption governs the way we determine the extension of 'morally right', it might seem that once we give up this assumption, there is no point in continuing to talk about 'morally right'. Of course, nothing in the definition of 'bound up' implies that the falsity of the assumption raises any challenge to the concept with which it is bound up, or to the practices involving the concept. However, whether this is the case in general or not, there is no doubt that challenging *either* of the reason-giving assumptions has important consequences to morality or the "moral point of view". The reason-giving assumption has a larger role to play in our moral life than being a convenient aid to help us figure out the extension of the concept 'morally right'. Morality would certainly lose its point if we no longer accepted that it gave us reasons for actions. To the extent that we are willing to agree that the nature of the

connection between those reasons for action and morality is a conceptual one, there seems to be no reason to stop short of full-blown internalism.

However, there is an important difference between weak internalism and full-blown internalism. If 'morally right' is indeed bound up with the reason-giving assumption, we can distinguish two uses of 'morally right', an agnostic and a committed one. When one uses the term 'morally right' in an agnostic way one does not take for granted the truth of the reason-giving assumption, even if one makes use of this assumption in determining what is morally right. On the other hand when one uses the term 'morally right' in the committed way, one has already accepted the reason-giving assumption. If full-blown internalism is true, one can use the term 'morally right' in its ordinary sense only if one remains committed to the truth of the reason-giving assumption. Given this fact, if full-blown internalism is right, a person who asks Glaucon's question has already answered it. But if weak internalism is true, we can use the term 'morally right' while remaining agnostic about the reason-giving assumption. One can thus ask Glaucon's question without thereby answering it; we *can* make sense of Glaucon's question without appealing to a Hare-substitute, or without any change in the meaning of 'morally right'. Since using the term 'morally right' does not commit us to the reason-giving assumption, the question "I see that this is morally right, but do I have a reason to do it?" can be seriously asked.

The internalist might protest that there is no reason why we should ever use the agnostic sense of 'morally right' once we accept the reason-giving assumption. Even if we can define this sense of 'morally right', there is nothing that it could express that could not be equally well expressed by the committed sense once we have accepted the assumption. Even if one grants that a weak internalist can ask Glaucon's question, one might object that on this view Glaucon's question just asks whether there is anything that is morally right at all, a question that a full-blown internalist can also raise.

Indeed the claim that we cannot challenge this assumption without challenging the moral point of view is compatible with weak internalism. It seems also right to say that if we were to concede to Glaucon that there is no reason to do what we are morally right to do, we should have conceded that there is a major flaw in the moral point of view, a flaw that might lead us to think that there is no place in our lives for moral requirements.³³ However, this does not mean that the question we asked was the same as the question whether there is anything that is morally right. First, a seemingly minor point: Glaucon's question is a more specific one than a general question about morality. There are many ways in which one can challenge morality, and this is a particular way. We are not asking whether some people feel pain or pleasure, or whether there are rational beings, or whether the concept of a person is metaphysically possible; all these questions might aim to cast doubt on the legitimacy of the moral point of view. Weak internalism allows us to "precipitate" a particular putative shortcoming of morality in a way that full-blown internalism does not allow for.

And the particular question that weak internalism allows to ask turns out to be an important one: Glaucon's question poses a challenge to the moral point of view that to some extent acknowledges its intelligibility, while calling its cogency into question. The person who asks Glaucon's question acknowledges that she can find her way inside the moral point of view, that she can see how things stand from there, (or at least that she understands the person who thinks that he can find his way inside this point of view³⁴), while she withholds her endorsement of it. The person who can raise questions from within the moral point of view, even questions that might undermine it, is in a different position from the person who sees that some actions get called "morally right" but cannot begin to understand why anyone would use such a label—someone for whom the morally point of view is utterly unintelligible.

The full-blown internalist may say that one might feel unmoved, or think that there is no reason to be moved, to perform any actions that fall under the extension of "morally right". Why wouldn't this position be the one that Glaucon asks Socrates how to avoid? And if this is so, isn't weak internalism once again just a notational variant of full-blown internalism? However, this could not work. If we accept full-blown internalism, the only answer to the above skeptic is to say: "Those actions that fall under the extension of 'morally right' are morally right, and thus, (as a matter of conceptual truth), one has a reason to do them".

It might be helpful to look at a couple of views that would be congenial to full-blown internalism, in order to see how weak internalism allows us to have a wider conceptual space. The first view could be called a "semi-reductionist" view. According to this view, the descriptive component of the concept "morally right," or of any other moral term, can be described in non-moral terms. In addition to this descriptive component, one may think that some moral concepts also have a non-descriptive component, which may or may not be capable of being reduced to non-moral terms. So fully reductionist views, views according to which the entire content of moral concepts can be captured in non-moral terms, would be a particular case of semi-reductionist view. I'll also take any view that there is no descriptive component to any moral concept (undoubtedly a quite implausible view, especially if we consider "thick concepts") as a limit case of a semi-reductionist view, so that the most extreme form of emotivism would count as a semi-reductionist view. The idea behind any semi-reductionist view is that the introduction of moral concepts to our vocabulary does not make any essential contribution to our "sorting" capacities. To the extent that those who have mastered the word "rude" have acquired the capacity to sort actions into those that are rude and those that are not³⁵ in a way that is invariant across those individuals who have mastered the word, it will be possible for someone who hasn't mastered any piece of moral vocabulary to acquire the same capacity. The rationale of a semi-reductionist view would be something like this: If human beings can sort the world according to a concept in roughly uniform ways so that a general agreement can be reached on the extension of this concept, it must be (or probably is) because the concept captures a genuine feature of reality to which human beings

are capable of responding. Since there are no irreducibly moral properties of the world, to the extent that a concept is tracking a genuine property of the world it is not an irreducible moral concept. So, if we find that there is a general agreement on the extension of a moral concept, we must be capable of forming a non-moral concept that can capture the same extension; in fact, if the agreement is non-accidental we must be capable of forming a concept that has the same extension in all possible worlds.

If one accepts a semi-reductionist view, there is no problem in taking one of the old sides of the dispute. One can use the non-moral concept to capture the extension of moral terms in all possible worlds, and Glaucon's question now can be formulated in terms of whether we have a reason to pursue those things. If one then thinks that the moral term contributes to the non-moral concept just by adding a positive answer to Glaucon's question, one will be an internalist. If one thinks that the main function of the moral term is to capture, perhaps in pre-scientific ways, what can be expressed in non-moral terms, one ought to be an externalist. The issue might seem like a verbal dispute, but it need not be. If one thinks, for instance, that the more general moral concepts have their content exhausted by the fact that their application marks the agent's endorsement of a certain type of action or a character trait, then one must think that externalism strips away all the content of those concepts. This person will also think that raising Glaucon's question would be raising an unintelligible question; it could be at best an imprecise manner of speaking or a form of doubting that the word has any application.

Indeed if I were to hazard a diagnosis of the debate, I would be inclined to say that, perhaps unwittingly, it is a semi-reductionist assumption that has governed the debate between internalism and externalism. But it is important to see that there is a different view, at the other extreme of the philosophical spectrum, that would also be congenial to the usual terms of the debate. If one thought that moral concepts could be competently applied only by those who *endorse* the moral point of view, then full-blown internalism would be a consequence of this view, and thus there would be no reason to advocate a weaker version of internalism.

However, if we reject both these assumptions, we are bound to think that the usual terms of the debate have not properly demarcated the conceptual space. That is, the acquisition of moral vocabulary might allow us to sort the world according to the assumptions, goals, and perhaps, the mode of understanding the world that are peculiar to the standpoint of morality, a standpoint that can occupied only by those who can appreciate the point of morality. If this is so, there is no guarantee that someone who has not acquired any moral vocabulary will be able to find a concept that matches the extension of any moral concept in the way required by a semi-reductionist view.³⁶ On the other hand, one need not assume that the standpoint of morality can be appreciated only by those who endorse it, and so one need not assume that the quick way to full-blown internalism outlined above is available.

If we reject these views that are congenial to full-blown internalism, we can see that weak internalism provides us with a plausible interpretation of Glaucon's question. The person who asks Glaucon's question, unlike the person who just asks whether anything is morally right, has placed herself among those who can understand the moral point of view. The possibility of understanding a point of view while withholding one's commitment to it—or while not endorsing this point of view—underscores our impression that there is an instance of Glaucon's question for which there is no Hare-substitute.³⁷

If someone points out the wide disagreement in matters of morality, and the constant revision of our opinions in this subject-matter, one is no doubt raising a challenge to morality; one is challenging our capacity to determine the *content* of morality. Glaucon's question has a different kind of concern in mind, a concern about the role morality plays in our lives; it is a challenge not to the internal organization of the system of morality, but to its point. If such challenges were raised only by skeptics, perhaps it would not be so important to keep them apart; they would be two items of the same general plan to undermine morality. However, the person who asks Glaucon's question—we might include Glaucon himself—is not necessarily a moral skeptic. One might seek a better understanding of the role that morality plays in our lives, not only for theoretical purposes but for practical purposes as well. I may see that, in certain cases, morality requires me to disregard the claims that the welfare of my siblings make on me in light of more impartial claims, or to leave behind projects that I have held dear. Here I might have no doubts about the *content* of morality; I am convinced that insofar as I continue to subscribe to the moral point of view I must leave the interests of my siblings behind. At this point, I might lose my grip on why I should not disregard instead the claims of morality. As Kant points out, morality has to win over the “seemingly plausible claims” of our competing inclinations,³⁸ and we will be more likely to succeed at this if we can improve our understanding of the nature of such claims. To this end, it is important to make intelligible a form of reasoning in which we withhold our commitment to the role the system plays in our lives without withholding our confidence in a certain mastery of the system. Of course reevaluating the role of the system might result in reevaluating its content, as a weak internalist view would lead one to expect. However, weak internalism makes it intelligible to take on these tasks one at a time.

A moral skeptic could also challenge the need of the reason-giving assumption in determining the extension of “morally right”. In fact, it could be the very nature of this challenge that it proposes a different way of determining the extension of “morally right”. We can interpret in this way, perhaps, the challenge that Thrasymachus, rather than Glaucon, raises in *The Republic*. According to Thrasymachus, “justice is simply the interest of the stronger”³⁹, and if we can generalize this claim to the concept “morally right”, the reason-giving assumption would not be bound up with the concept “morally right”. It is part of the argumentative burden of the weak internalist to be capable of dismantling such proposals, to show that Thrasymachus cannot capture this way the concept “morally right”.⁴⁰

However, one could also interpret Thrasymachus as providing something closer to a genealogy rather than a definition of the concept. In this case, Thrasymachus would be claiming that the system of morality has as its function to service the interest of the stronger. This claim could presumably cast doubt upon the reason-giving assumption. If this is true, we can see Thrasymachus as providing a perhaps less plausible version of Glaucon's challenge—a different, perhaps less convincing, reason to doubt that we have reason to do that which is morally right. In this cases, Thrasymachus's skepticism, just as Glaucon's doubt, can be best understood within the framework provided by weak internalism.

V

But is the concept 'morally right' in fact bound up with the reason-giving assumption? Is it true that determining what counts as morally right depends on this assumption? One way to try to establish this claim is to rely on intuitions about linguistic awkwardness. Smith presents as evidence for (full-blown) internalism the following case:

Suppose we debate the pros and cons of giving to famine relief and you convince me that I should give. However when the occasion arises for me to hand over my money I say 'But wait! I know I *should* give to famine relief. But you haven't convinced me that I have any *reason* to do so!' ⁴¹

It does indeed seem that we would be puzzled by what this person is saying. However, this line of defense is somewhat problematic. It is far from clear that the puzzlement results from a conceptual connection between calling something 'morally right' and having a reason to do it.⁴² It might be a pragmatic implicature of endorsing the claim in this context; after all there would be no reason for someone to engage in a long dispute about the morality of giving money to famine relief if she did not think that there were reasons to do what is morally right. Moreover, even the externalist thinks that we have reasons to do what is morally right, even if those reasons are not part of the concept 'morally right'. Since the externalist typically thinks that there is a *substantive* connection between that which is morally right and that which we have a reason to do, she can also explain why we express puzzlement at those claims. For even if this is not a conceptual connection, it still might be a substantive connection that we generally take for granted, and so, we shall naturally be surprised to find it denied.

However, our discussion of weak internalism in the case of the concept 'Kosher' should give us better reason to accept the truth of weak internalism in ethics. It is hard to make sense of any form of ethical reasoning if we do not take the reason-giving assumption to be true. Let us take, for instance, as a form of reasoning the requirement that we can universalize moral demands. It seems legitimate in a moral argument to ask, "what if everyone did X?". However, if the concept of morally right were not bound up with the reason-giving assumption, it

would seem appropriate to respond: "Indeed that would be awful. But I was suggesting only that it was morally right. I was not suggesting that anyone should do it." The same holds when someone tries to show that something is morally wrong by asking: "How would you like it if someone did this to you?" It would not be a legitimate rebuttal to answer: "By saying that it was morally right to do X, I was not suggesting that anyone would have a reason to do it."

Of course, the externalist could raise the same objection again. Since the externalist claims that there are substantive connections where the full-blown and weak internalists see conceptual ones, she might claim that we find these rebuttals odd because they deny a widely accepted substantive connection. The distinction between conceptual and substantive connections is no doubt a murky issue, but if we can make any sense of it, we must have some means of distinguishing between the two. And here, it might be worth invoking an old criterion; if the connection between 'morally right' and the reason-giving assumption is not a conceptual one in the way suggested by weak internalism, we must be capable at least to conceive certain circumstances in which this connection does not obtain, and in which the above rebuttals would be legitimate.⁴³ Since ordinary ethical reasoning (and even more so *philosophical* ethical reasoning) typically contemplates imagined as well as real situations without ever giving up these forms of reasoning, and since these forms of reasoning seem to be at the core of all our thinking about morality, it seems unlikely that we can conceive of such circumstances. At the very least, it seems not only commodious but also appropriate to shift the burden of the proof. The externalist needs to show how we can conceive of circumstances in which the above rebuttals can be legitimately given without changing the subject.⁴⁴

To the extent that externalism claims that there is no conceptual connection between the reason-giving assumption and claims of what is morally right, we lose the sense of how ethical reasoning could possibly be adequate. The externalist can make sense of Glaucon's question only by failing to make sense of moral reasoning in general. However, one may say, many externalists are also naturalists. As a naturalist, one might think that the categories of morality have to be justified in the same terms that we justify categories in the natural sciences; that is, by their role in the causal explanation of events in the natural world. This is no doubt a somewhat simplistic characterization of naturalism, but even if naturalism thus characterized were a correct view it would not undermine the above argument in favor of internalism. For if ethical naturalists do not want to change the subject matter, they must be giving an account of the ordinary concept of morality. Even if a naturalist would gladly accept some revisions in our ordinary concept, she would not want to depart from it so much that one could no longer recognize the original concept in her account.⁴⁵ In order to provide such an account, we must first identify the commitments involved in our ordinary concept of morality. And the above considerations suggest that certain forms of reasoning that depend on the reason-giving assumption are constitutive of our concept of morality. Thus, to the extent one claims to be providing an adequate account of

the ordinary concept 'morally right', one must be purporting to provide an account of a concept that is bound up with the reason-giving assumption.

A certain brand of externalism claims that we need to separate two things: the properties that make something morally right and the reasons we have to perform actions that exemplify this property. Externalists take the possibility of Glaucon's question to be evidence for this view. If Glaucon's question makes sense, it must be possible in principle to determine these properties in a way that is independent of any considerations about what we have reason to do or not to do. It is my contention, however, that as long as Glaucon can determine what counts as morally right, he must still rely on the assumption that if X is morally right then there is a reason to X. Thus the possibility of Glaucon's question does not give evidence that the property of being morally right can be identified with a natural property such that we can determine whether something exemplifies this property or not independently of considerations of whether it can provide reasons to act or motivation. In other words, if I am right, a naturalistic account of moral properties will face the same difficulties that it would face if full-blown internalism were true. Of course, this is not a fatal blow to this kind of moral realism. One might still want to say that a more radical revision is needed in our moral concepts, or advocate an error theory with respect to ordinary moral discourse together with a substitute notion of morality that can answer the constraints imposed by naturalism. However, one must give up the claim that the intelligibility of Glaucon's question stands witness to the fact that the concepts of ordinary morality already embody the form of externalism demanded by this kind of moral realism.⁴⁶

Notes

¹To avoid needless repetition, I will often omit the qualifier '*prima-facie*'. I find the expression 'morally right' more neutral with regard to different moral theories (and less cumbersome) than other expressions that would not require this qualification, such as 'morally required' and 'morally obligatory'. For a precedent for choosing 'morally right' over these alternatives, see W. D. Ross, *The Good and the Right* (Indianapolis: Hackett Publishing Co., 1988), p. 4.

²Cp. David Brink: "Broadly speaking, internalism is the view that there is an internal or conceptual connection between moral considerations and actions or the source of actions" in *Moral Realism and the Foundations of Ethics* (New York: Cambridge University Press, 1989), p. 38. Brink also gives a long list of possible ways of clarifying this connection. See *Moral Realism*, pp. 37–43. Brink correctly points out that internalism can be defined in terms of motives rather than reasons. I will not be concerned with this form of internalism below.

³Proponents of ethical internalism include Richard Hare in *Language of Morals* (New York: Oxford University Press, 1952), J. L. Mackie in *Ethics: Inventing Right and Wrong* (New York: Penguin Books, 1977), Michael Smith in *The Moral Problem* (Oxford: Blackwell Publisher, 1994). Proponents of ethical externalism include William Frankena in "Obligation and Value in the Metaphysics of G. E. Moore" in Paul Arthur Schilpp (ed.), *The Philosophy of G. E. Moore* (Evanston, IL.: Northwestern University Press, 1942), Brink in *Moral Realism and the Foundations of Ethics* and Peter Railton, "What the Non-Cognitivist Helps us to see the Naturalist Must Help Us to Explain" in John Haldane and Crispin Wright (eds.), *Reality, Representation and Projection* (New York: Oxford University Press, 1993).

⁴For the sake of (near) precision: We could define an empty question as a question such that some of its presuppositions obviously imply a direct answer to the question. For precise definitions of ‘presupposition’ and ‘direct answers’ to a question, see Nuel Belnap and Thomas Steel, *The Logic of Questions and Answers* (New Haven: Yale University Press, 1976).

⁵This is different, however, from Darwall’s own distinction between judgment and existence internalism since we are dealing with different *relata*. Darwall is concerned, in this context, with the relation between reasons to act (or morality) on the one hand and motives on the other hand, rather than between on the one hand the concept “morally right” and reasons to act on the other hand. See *Impartial Reason* (Ithaca, NY: Cornell University Press, 1983), pp. 51ff. Brink makes a similar distinction between appraiser and agent internalism. See *Moral Realism*, p. 40.

⁶Strictly speaking, there is a third option. That is to argue that no matter what ends one may have, doing what is morally right is the best means to further those ends. For a defense of this option, see David Gauthier, *Morals by Agreement* (New York: Oxford University Press, 1986).

⁷This does not rule out the possibility that one or both of them are *in fact* (or even necessarily, but not conceptually) true; this possibility is left open by the paper.

⁸Of course, the same goes, *mutatis mutandis*, for existence internalism and externalism.

⁹I am using “instrumental conception of rationality” because this is the notion used by Peter Railton in “Moral Realism”, *Philosophical Review* 95, 1986, pp. 163–207. The same would be true, however, of maximizing conceptions of rationality. For the distinction between these two conceptions of rationality see David Gauthier, “Reason and Maximization” in *Moral Dealing* (Ithaca: Cornell University Press, 1990).

¹⁰I think J. L. Mackie comes closer than Hume to provide such an argument in *Ethics*, ch. 1. The argument is attributed to Hume in essentially this form by Railton in “Moral Realism”, pp. 167–8.

¹¹See, for instance, Railton, “Moral Realism”, pp. 167ff, and Brink, *Moral Realism*, pp. 43ff, for defenses of this possibility.

¹²Although I will focus on the form of moral realism that advocates that moral properties are natural properties, what I am saying here would also apply to those who identify moral properties with *non-natural* properties, such as G. E. Moore. Moore himself thinks that he is committed to an externalist view of morality. But since Moore’s view has largely fallen out of favor, it seems worth ignoring it for the sake of stylistic elegance.

¹³The reason for this qualification should be clear below at pp. 118–19.

¹⁴Similarly, Simon Blackburn concedes that the reason-giving₃ assumption is not true universally (in his words, “externalists can win individual battles”), but claims that this is compatible with assuming that internalism is the correct position (in his words, “internalists win the war”), since, according to Blackburn, the cases in which the externalists win the battle are necessarily “parasitic”. See his *Ruling Passions* (New York: Oxford University Press, 1998), p. 61. Although Blackburn does not give us much reason to accept that these cases indeed need to be parasitic, I think he is right that if they are, then the internalist position is essentially right, since it is still true that there is a conceptual connection between ‘morally right’ and the reason-giving assumption.

¹⁵I will return to this point at the end of the paper.

¹⁶See Brink’s *Moral Realism*, pp. 57–9, for this kind of criticism of internalism.

¹⁷*Republic*, trans. by G. M. A. Grube (Indianapolis: Hackett Publishing Co.: 1992), Book II. Nothing in this paper depends on the question whether this is a correct interpretation of the *Republic*.

¹⁸*Language of Morals*, p. 164.

¹⁹Since Hare suggests that different instances of this question have different meanings (See *Language of Morals*, p.167), it is important that the definition refers to the meaning of *instances* of the question, instead of to the meaning of the question itself.

²⁰Of course, ‘that which is generally called “morally right”’ mentions, but does use a moral term. Note also that replacing ‘morally right’ by any expression that means the same as ‘morally right’ (even if, assuming that this kind of reduction is possible, the same content is expressed in, say, the vocabulary of fundamental physics) would not enable the internalist to raise a non-empty version of Glaucon’s question. If the internalist is right, any expression E that has the same mean-

ing as ‘morally right’ must be such that if one accepts that X is E then one has to accept that there is a reason to do X .

²¹Compare Alfred Mele’s definition of “MR belief*” in “Moral Cognitivism and Listlessness”, *Ethics* 106, 1996.

²²*The Moral Problem*, pp. 68–71.

²³Smith cites Christopher Peacocke as an example of such a theorist. See Peacocke’s *Sense and Content* (New York: Oxford University Press, 1985).

²⁴This would not be true if we could somehow import the content of the blind person’s expression ‘blue’ into our language. I am assuming that, since this concept would depend on a kind of experience we lack (just as our color concepts depend on kinds of experience the blind person lacks), it could not be imported into our language.

²⁵*Moral Realism*, p. 59, emphasis mine.

²⁶I will reexamine the nature of the evidence for internalism in section V.

²⁷I am assuming that the “Kosher” stamp that we find on many food items is not constitutive of what counts as “Kosher”, but it is rather only a reliable guide to whether something is Kosher or not, much like FDA approval, which might be quite reliable, but nonetheless not an infallible guide to which food is healthy or not. We can think of those responsible for the stamp as those whose assessment is trusted by the Jewish community, rather than as those who decide what *counts* as Kosher. The truth of the matter is more complicated than this, but this is not very far from the truth, since presumably an orthodox Jew who finds something that is obviously a slice of pork in a container with the Kosher stamp should count this as an oversight and not just think that he has been fortunate enough to run into the only slice of Kosher pork in the world.

²⁸Or at least X will not be important for her in any standard way.

²⁹Or of any other tradition or complex mode of reasoning.

³⁰Also, it might seem that the fact that ‘Kosher’ is bound up with a certain assumption plays a role in determining its extension only in borderline cases. Although I do not think that this is true even of the concept of Kosher, it is certainly not necessarily true for every case in which a concept is bound up with an assumption.

³¹Of course there are many trivial ways of constructing an axiomatic system in which we can define a predicate ‘Kosher’ such that all true sentences of the form ‘ X is Kosher’ would follow from it (such as the system that contain only the axiom “everything that can correctly be called ‘Kosher’ is Kosher, and nothing else is Kosher”). The important point is that we cannot take any of the Rabbis to be making implicit or explicit use of any such axiomatic system in settling whether something is Kosher or not.

³²The reason I leave this point to the end should become clear, I hope, by then.

³³This is in fact a welcome consequence of the view, since there is no doubt that someone who is posing Glaucon’s question is contemplating the possibility that morality might have no role to play in our lives.

³⁴Cf. Cora Diamond’s interpretation of Wittgenstein’s claim that the propositions of the *Tractatus* are nonsense in “Ethics, Imagination and the Method of Wittgenstein’s *Tractatus*” in *Bilder der Philosophie*, Wiener Reihe 5 (Vienna, 1991).

³⁵I am ignoring some complications. Of course, mastering a concept X does not involve a flawless competence in sorting those things that are X from those that are not (in case of some natural kinds it is not clear that one need to have any kind of competence like that), and mastering a concept involves inferential capacities that are not easily reduced to classificatory capacities. I hope it is clear that this simplification is innocuous.

³⁶For a similar view with respect to thick concepts, see Williams’s famous discussion of thick concepts in *Ethics and the Limits of Philosophy* (Cambridge, Mass: Harvard University Press, 1985). For a similar view with respect to values in general, see John McDowell, “Values and Secondary Qualities” in his *Mind, Values and Reality* (Cambridge, Mass.: Harvard University Press, 1998).

³⁷Someone might think that it is not possible to understand the moral point of view *fully* without endorsing it. If this view were right, then, it would follow from weak internalism that anyone who

raises Glaucon's question seriously does not *fully* understand the moral point of view (but this person would still need to have *some* understanding of the moral point of view to be capable of raising Glaucon's question meaningfully).

³⁸*Grounding for the Metaphysics of Morals*, trans. by James Ellington, (Indianapolis: Hackett Publishing Company, 1980), p. 16 (Ak. 405).

³⁹*Republic*, 338c

⁴⁰This is not to say that Thrasymachus does not *use* the concept in the same way we do. He might be giving just a bad *account* of the concept "morally right", and the fact that he uses it in the same way we do might be important to show that his analysis is inadequate. More plausibly, however, Thrasymachus might be trying to show that people who use this concept are often prey to a confusion, and that despite the fact that they would not recognize the concept "morally right" in Thrasymachus's account, this is the only way one could *meaningfully* use it.

⁴¹*The Moral Problem*, p. 60.

⁴²Indeed the fact that the unqualified word 'should' is being used might be part of the reason why this seems such an awkward thing to say. After all the word 'should' is used to denote having reasons for actions in general ("I should consume more dietary fiber").

⁴³Note that externalists such as Brink use exactly this criterion to establish the claim that there is no conceptual connection between 'morally right' and the reason-giving assumption. The claim is supposed to be established by the fact that we can meaningfully ask Glaucon's question.

⁴⁴Note that the reply I give here is not available to Smith. It is all too easy to imagine circumstances in which it would be appropriate to say: "I accept that it is morally right to give to famine relief. But you haven't convinced me that I have any *reason* to do so." We just need to imagine that in the middle of our conversation about the morality of giving to famine relief, I express that I have been having doubts about morality, and I am no longer sure one has always a reason to do what is morally right. We then continue our discussion, and you convince me that indeed it is morally right to give to famine relief.

⁴⁵Railton, for instance, acknowledges that "revisionism may reach a point where it becomes more perspicacious to say that a concept has been abandoned, rather than revised. No sharp line separates tolerable revisionism and outright abandonment, but if our naturalist wishes to make his case compelling, he must show that his account (...) is a rather clear case of tolerable revision, at worst." In "Naturalism and Prescriptivity", *Social Philosophy and Policy*, vol. 7 (1), 1989, p. 159.

⁴⁶I would like to thank James Doyle, Russell Goodman, Hans Lottenbach, Jennifer Nagel, Amy Schmitter, Fred Schueler, John Taber, Aladdin Yaqub, and two anonymous referees for *Noûs* for extremely helpful comments on earlier drafts of the paper.