


This article was downloaded by: [University of Oxford]

On: 28 July 2009

Access details: Access Details: [subscription number 773573598]

Publisher Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Philosophical Explorations

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713706422>

Humean agent-neutral reasons?

Daan Evers^a

^a Faculty of Philosophy, University of Oxford, Oxford, UK

Online Publication Date: 01 March 2009

To cite this Article Evers, Daan(2009)'Humean agent-neutral reasons?',Philosophical Explorations,12:1,55 — 67

To link to this Article: DOI: 10.1080/13869790802635614

URL: <http://dx.doi.org/10.1080/13869790802635614>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Humean agent-neutral reasons?

Daan Evers*

Faculty of Philosophy, University of Oxford, Oxford, UK

In his recent book *Slaves of the Passions*, Mark Schroeder defends a Humean account of practical reasons (*hypotheticalism*). He argues that it is compatible with ‘genuinely agent-neutral reasons’. These are reasons that any agent whatsoever has. According to Schroeder, they may well include moral reasons. Furthermore, he proposes a novel account of a reason’s weight, which is supposed to vindicate the claim that agent-neutral reasons (*if* they exist), would be weighty irrespective of anyone’s desires. If the argument is successful, it could help avoid an error-theory of moral language. I argue that it isn’t, and that we should reject a Humean approach to reasons.

Keywords: Mark Schroeder; hypotheticalism; practical reasons; weighing reasons; instrumentalism; moral judgment; error-theory

1. Introduction

Consider the view that whether *X* has a reason to *A* depends on *X*’s having some desire that is furthered by his *A*-ing (Humeanism about practical reasons). This type of view has great appeal to many, but meets with equal opposition. The appeal is explained by two considerations: first, the theory seems to explain certain simple cases well. Mark Schroeder gives the following example: if Ronnie loves to dance, but Bradley doesn’t, then the fact that there is dancing at the party is a reason for Ronnie to go, but not for Bradley (2007, p. 1). Humeanism predicts this: the dancing at the party is a reason for Ronnie to go, because his going there enables him to dance, which he desires to do. Since Bradley lacks the desire, he does not have this reason. Humeanism, then, seems to get something right. The second consideration which explains its appeal is metaphysical: if reasons for action are a function of an agent’s desires, there is no need to postulate obscure or unexplained normative entities (see, e.g., Harman, 1985). So Humeanism has some clear advantages. But it also seems to have important drawbacks.

One of these is that it makes all reasons hypothetical (as opposed to categorical): Ronnie’s reason to go to the party is conditional on his desire to dance. But are all reasons like this? Many people think that arbitrary torture is wrong, and that its wrongness is independent of the desires of an agent (any agent). So it seems that ‘*X* ought not to torture arbitrarily’ is true, irrespective of anyone’s (including *X*’s) desires. But, presumably, ‘*X* ought not to *A*’ entails that *X* has (most) reason not to *A*. However, if moral statements entail statements about reasons, then the truth of the first depends on the truth of the latter. So if moral statements are categorically true (or false), then the corresponding reason statements ought to

*Email: daan.evers@jesus.ox.ac.uk

be as well. But if Humeanism is correct, then whether X has reason to A depends on whether A -ing furthers his desires. And so whether X ought morally to A depends on whether A -ing furthers his desires. And that undermines the categorical nature of moral truths.

Quite a few philosophers believe that the categorical nature of morals is built into moral language (e.g. Joyce, 2001; Kant, 1785/1998; Mackie, 1977). According to these philosophers, it is a matter of conceptual fact that moral truths (and therefore moral reasons) are independent of desires. If so, and if Humeanism is true, it seems to follow that all (positive and nonanalytic) moral statements are false. For Humeanism seems to entail that there *are* no reasons which are independent of desires. So the truth of Humeanism seems to entail an error-theory of moral language (as defended by Mackie, 1977 and Joyce, 2001).

But even if Humeanism were able to account for categorical reasons, this may not be enough to save moral language from systematic falsehood. Richard Joyce argues that moral reasons are (as a matter of conceptual fact) not only independent of desires, but also *weighty* independent of desires (2001, chap. 2; a similar view is defended by Richard Hare concerning moral judgments in 1981, chap. 3). Moral reasons are (at least in typical cases) *overriding* or stronger than other types of reasons. But if reasons depend on desires, their weight is naturally thought of as a function of them too. This leads to what Mark Schroeder calls *proportionalism*: ‘the thesis that when a reason is explained by a desire, as in Ronnie’s case, its weight varies in proportion to the strength of that desire, and to how well the action promotes that desire’ (2007, p. 98). Proportionalism thus entails that whether X ’s moral reason to A is strong or weak depends on the strength of his desire for something which is furthered by his A -ing. If X ’s desire is weaker than other desires of his (which are not furthered by his A -ing), his moral reason will be weak as well. But, intuitively, X has strong moral reason to stop torturing the cat even if he has *no* desire which is furthered by that action.

In the light of the above, Humeans should welcome Schroeder’s book *Slaves of the Passions*. For Schroeder argues that his version of Humeanism is compatible with the existence of ‘genuinely agent-neutral reasons’ (including moral ones). According to *hypotheticalism*, there are reasons that anyone with any desires at all must have. Furthermore, Schroeder argues that if there *are* such reasons, they will be *weighty* irrespective of anyone’s desires. So hypotheticalism promises to save moral language from systematic falsehood. This would be a significant achievement. However, impressive and original as Schroeder’s arguments may be, I think he fails to establish either that hypotheticalism implies that moral reasons are really agent-neutral, or that they would be weighty irrespective of desires.

2. Hypotheticalism

Schroeder distinguishes several different aspects of the (supposed) objectivity of moral reasons:

- (1) *Agent-neutrality*: (K) ‘The fact that Katie needs help is a reason to help her’ does not explicitly refer to anyone *for whom* it is a reason. It is agent-neutral *in form*.

But Schroeder believes all reason statements do, albeit sometimes implicitly, refer to a class of people, as follows:

Agent-neutral For R to be a reason to do A is for R to be an agent-relational reason for all of [us] to do A . (2007, p. 18)

However, in the case of moral reasons, ‘us’ includes *everyone*, which is the second objectivity feature that Schroeder distinguishes:

- (2) *Universality*: in cases like (K), ‘us’ includes everyone.¹

The third feature that Schroeder distinguishes is related to ‘the fact that the universality of the reason to help Katie does not seem to be contingent – for example, on the fact that the only people around happen to value Katie’s welfare’ (2007, p. 105). He then distinguishes two modal claims:

- (3) *Weak modal status*: moral reasons do not depend on the contingent desires of anyone in particular.

This means that (2) should be read as saying: ‘in cases like (K), “us” necessarily includes everyone’. But (3) is still compatible with moral reasons depending on certain specific desires, namely necessary ones. However, Schroeder claims that something stronger seems the case for moral reasons, namely:

- (4) *Strong modal status*: ‘for any desire, an agent would have a reason to help Katie even if he did not have that desire’ (2007, p. 106).

This rules out that moral reasons might depend on specific necessary desires.

According to Schroeder, it is not hard to account for the agent-neutrality of reason ascriptions on a Humean understanding of reasons. It suffices that we understand agent-neutral ascriptions as elliptical for agent-relational ones, where the scope of ‘us’ is contextually determined and sometimes (for example in the case of moral reasons) includes everyone. If ‘everyone’ is thought of as everyone possible, then reason ascriptions will have the weak modal status.

The above story of Schroeder’s is of course semantical: it is an explanation of how reason statements could have the kind of commitments often thought to be features of moral discourse. But the explanation isn’t particularly Humean, except perhaps in its claim that all reason statements are implicitly agent-relational. However, not even this feature need be Humean. For reasons to do things are reasons for agents (possibly all, or even necessarily all). It makes little sense to speak of reasons that are there, but that are reasons for nobody.

I am happy to grant that reason statements can be understood in the way Schroeder suggests. So I’ll assume that all reason statements, including superficially agent-neutral ones, are really agent-relational in disguise. The scope of ‘us’ may be contextually determined. The question is whether moral reason statements turn out to be false *en masse* when ‘us’ is construed so as to include all possible human (and perhaps even all rational) beings. In other words, we want to know if there *are* any (moral) reasons that everybody has. This is where Schroeder shows great philosophical ingenuity. He argues that there may well be.

Step one in his argument is to note that it is compatible with the view that reasons depend on desires that reasons don’t depend on any particular desire. It needn’t be the case that if *X* lacks desire D_1 , he doesn’t have a reason to *A*. For the reason to *A* may also be grounded in D_2 as well, which *X* *does* have. If a reason *R* to *A* can be grounded in several different desires, then the following is true: there is no specific desire D_n such that if *X* lacked it, it would follow that *R* isn’t a reason for him to *A*.

I agree that most reasons can be grounded in several desires. An example of Schroeder's serves to illustrate the point:

Susan wants some coffee. So the fact that there is coffee in the lounge is a reason for her to go there. But perhaps this is doubly determined. For example, perhaps philosophers tend to congregate and talk shop when there is coffee. And perhaps Susan wants to talk shop about some idea she's recently had. This should also explain why the fact that there is coffee in the lounge is a reason for Susan to go there. If this is so, then this is a reason for Susan twice over. The fact that there is coffee in the lounge would be a reason for Susan to go there even if she didn't want a cup of coffee, and it would be a reason for her to go there even if she didn't want to talk shop. So there is no single desire on which it depends. (2007, pp. 108–109)

Of course, this is still not enough to establish that there are any reasons that anyone at all must have. For even if there is not one single desire on which Susan's reason to go to the lounge depends, it depends on a small number which she happens to have. People who have neither a desire for coffee, nor for talking philosophy, don't have a reason to go to the lounge in Susan's building. Few philosophers hold that the objectivity of moral reasons would be sufficiently guaranteed if it turned out that they depend on [either a desire for p , or a desire for r , or ...], where this is a relatively short disjunction.

What Schroeder proposes is that some reasons (amongst which moral ones) depend on almost any desire:

Hypotheticalism's favored proposal for how there could be genuinely agent-neutral reasons is therefore that genuinely agent-neutral reasons are *massively* overdetermined. They are reasons for anyone, no matter what she desires, simply because they can be explained by any (or virtually any) possible desire. (2007, p. 109)

Schroeder, then, believes that if there are reasons that are explained by any possible desire, they will be reasons for anyone, no matter what he desires, as long as there is *something* he desires. So genuinely agent-neutral reasons depend on desires, but in a very minimal sense. For, as Schroeder claims, it isn't particularly plausible that beings without desires at all are agents in the first place. And it *is* plausible that agenthood is a necessary condition on having reasons.

Next, Schroeder gives the following definition of a reason:

Reason For R to be a reason for X to do A is for there to be some p such that X has a desire whose object is p , and the truth of R is part of what explains why X 's doing A promotes p . (2007, p. 59)

Now, the basic form that Schroeder's explanation takes (of the idea that certain reasons may be had in virtue of any possible desire) is illustrated by the example of a reason to believe an arbitrary proposition only if it is true. Here is how the explanation goes:

Being in error about [some arbitrary proposition] might lead to being in error about other things, such that being in error about them might lead to being in error about other things, and so on until something might lead to Mary having trouble getting new shoes. If this is right, then for any proposition, Mary's desire to get a new pair of shoes will serve to explain why there is a reason for Mary to believe it only if it is true. (2007, p. 114)

So, the basic idea is that for any proposition, there is a chance that *if* it is true, believing it could come to bear on Mary's ability or success at buying shoes (being in error *might* lead to

having trouble buying shoes). Therefore Mary has a reason to believe any arbitrary proposition only if it is true in virtue of her desire to buy new shoes.

I assume that the proposition whose truth plays a role in an explanation of why doing *A* promotes *p* has to be true *in the actual world*. That the fairies exist cannot be a reason to do something if the fairies don't exist. But it would be odd if all that the proposition has to do is play a role in an explanation of why doing *A* would promote *p* in some *possible* world (for example, a world where the laws of nature are completely different). Doing *A* has to promote *p* in the actual world in order for you to have reason to do *A* in the actual world. So I will also assume that the proposition that is true in the actual world has to play a role in an explanation of why doing *A* would promote *p* in the actual world.

Partly for independent reasons, Schroeder takes what he thinks is a very permissive stance on what it is for an action to *promote* an object of desire:

X's doing *A* promotes *p* just in case it increases the likelihood of *p* relative to some baseline. And the baseline, I suggest, is fixed by the likelihood of *p* conditional on *X*'s doing nothing – conditional on the status quo. (2007, p. 113)

The idea behind making the likelihood of *p* relative to doing nothing is to ensure that it is easy to generate reasons. As long as it is *slightly* more likely that *p* will occur if *X* *A*-s than if he does nothing, *X* has a reason to *A*. And the easier it is to generate reasons, the more likely it is that there will be genuinely agent-neutral ones.

3. Humean agent-neutral reasons?

So does Schroeder succeed in showing that there are reasons (including moral ones) which anyone with any desires at all is bound to have? Before I examine this, we need to ask whether the success of Schroeder's argument would actually save moral language from an error-theory. One could take issue with Schroeder's characterization of the kind of objectivity that philosophers like Joyce believe accrues to moral reasons (or would accrue to them, if such reasons would exist). Combining the weak and strong modal statuses, one could say that Schroeder thinks the relevant type of objectivity is this:

Specific desire independence: for any specific desire *D* (whether contingent or necessary), it is not the case that an agent would lack a moral reason (e.g. to help Katie) if he did not have *D*.

But it is not at all unlikely that what philosophers like Kant, Mackie and Joyce have in mind is this (or a principle like it):

Wholesale desire independence: for any moral reason *R* of an agent *X*, it is not the case that *X*'s having *R* at *t* depends on the content of any specific or several of his desires at *t* or any other time.

Wholesale desire independence can be true even if having some desires is a necessary condition on being an agent (and thus an entity capable of having reasons in the first place). What it entails is that *X*'s now having a reason to stop torturing the cat is nothing to do with whether his actual desires are (likely) to be served by it (it does *not* entail that *X* would have the reason even if he had no desires whatsoever). But hypotheticalism does not allow this. For it says that *R* is a reason for *X* to *A* if and only if there is some *p* such that *X* has a desire whose object is *p*, and the truth of *R* is part of what explains

why X 's doing A promotes p . So X 's having R at t depends on the fact that he desires something specific, namely p .

Many philosophers would still find this objectionable. No matter how easy it is to generate moral reasons from desires, the very fact that they ultimately depend on their (foreseeable) instrumental role in promoting our desires is what drives these philosophers away from Humeanism (see, e.g., chap. 3 of Derek Parfit's unpublished manuscript *On What Matters*, 2008). But that is not what matters now. What matters now is that if *Wholesale desire independence* is built into the concept of a moral reason, and hypotheticalism does not allow it, then even if everything else Schroeder says is true, an error-theory has still not been averted. But, unfortunately for him, he also fails to avert it on his own terms (those described at the start of Section 2).

We have seen that Schroeder's strategy in arguing for agent-neutral reasons is to have a very permissive criterion for what it is for an action to promote an object of desire. He says that X 's doing A promotes p just in case it increases the likelihood of p relative to doing nothing or the status quo. But this definition may not generate reasons as easily as Schroeder hopes. First, we need to get clear on what 'doing nothing' is supposed to be. It can hardly be sitting absolutely still. For sitting absolutely still is also doing something for which one may have reasons. But if sitting absolutely still is the baseline relative to which *all* reasons are determined, then one can never have any reasons to sit absolutely still. After all, whatever the character of p , the likelihood of realising it by sitting absolutely still cannot be higher than the likelihood of realising it by sitting absolutely still.

Another candidate for 'doing nothing' is *anything that constitutes not doing A*, which can be many things. For example, if A is giving Katie 100 pounds, then, on this interpretation, I'm not doing A if I give her 99 or 101. But it may well be that relative to those baselines, the likelihood of my realising p is not in fact raised at all. So perhaps this is not what Schroeder had in mind either.

Since Schroeder talks about 'the status quo', it may seem that 'doing nothing' is something *specific* that constitutes not doing A : something like *continuing to do what one was already doing*.² For instance, if what one is currently doing is staring out of the window, then 'doing nothing' is not giving Katie 99 pounds, but continuing to stare out of the window. But this interpretation creates the same problem as the first: surely one can have reasons to continue doing what one was already doing. But whatever the objects of one's desires, the likelihood of realising them by continuing to C cannot be higher compared to the likelihood of realising them by doing the very same.

So it is not exactly clear what 'doing nothing' amounts to. But this question is important. For Schroeder needs it to be easy to raise the likelihood of p . But this may be much harder relative to some baselines than others. And that is problematic for his genuinely agent-neutral reasons. In fact, not even his most convincing case seems to work: it is false that any desire whatsoever can ground a reason to believe an arbitrary proposition only if it is true.

Consider whether Mary has a reason to believe a true, but highly abstract proposition about metaphysics in virtue of her desire to buy new shoes. Suppose there is some non-zero probability that believing this proposition affects her success at buying shoes. There is presumably also a non-zero probability that not believing it will *not* affect her success. Suppose these probabilities are equal (which is a charitable assumption; the probability that not believing the proposition will *not* affect Mary's success at buying shoes may well be greater than the probability that believing the proposition *will* affect it). If so, and if we fix the baseline ('doing nothing') at *not believing the proposition*, then relative to it, the likelihood of p (buying new shoes) is not in fact raised. This

means that relative to not believing the proposition, Mary does *not* have a reason to believe the proposition (in virtue of her desire to buy new shoes). But then relative to what *does* she have this reason, since not believing the proposition and believing the proposition cover all the options?

So it is much harder to find any genuinely agent-neutral reasons than Schroeder thinks. This also holds for moral reasons. Consider the fact that Katie needs help. Clearly, logical space contains many scenarios in which helping Katie promotes any arbitrary goal of mine. But that is not enough. The fact that Katie needs help is a reason for me to help her in virtue of some desire of mine only if there is some (actually) true proposition such that it plays a role in an explanation of why helping her raises the probability of realising that desire in the actual world. As before, it all depends on the baseline. Most desires (for ice creams, seeing my parents, helping someone other than Katie, etc.) are so far removed from anything that might be achieved by helping Katie, that the likelihood of attaining their objects is not going to be higher if I do help her compared to doing something else instead. But whatever 'doing nothing' is supposed to mean exactly, it will presumably at least involve doing something else instead.

Lastly, we may wonder what makes a baseline the correct one. If there are no objective facts about which baselines are correct, then there are no objective facts about what is and is not an agent-neutral reason. Or at least, there will only be objective relational facts. But the fact that *X* has a reason to stop torturing the cat *relative to* B_1 will not satisfy the moral objectivist. For there is likely to be some baseline relative to which *X* does *not* have a reason to stop torturing the cat. And if so, it is indeterminate (and thus not true or false *simpliciter*) whether *X* has a reason to stop torturing the cat.

4. Weighty agent-neutral reasons?

We have already seen reasons to doubt Schroeder's claim that everyone has moral reasons in virtue of any possible desire (or virtually any possible desire). This means that if the universality of moral reasons is implicit in moral language (as specified in (2) and (3) in Section 2), we may have to embrace an error-theory (or some kind of revisionary semantics). At least this would be so if Humeanism is the correct account of practical reasons.

But even if Schroeder did succeed in explaining why everyone has moral reasons, an error-theory would still be lurking if those reasons turn out to be highly variable in strength. Many people feel that moral reasons are strong reasons that often trump other kinds of reasons. Schroeder argues that, *if* it can be shown that anyone has moral reasons in virtue of (almost) any possible desire, then those reasons will be weighty for all people.

Step one in the argument is the rejection of proportionalism: the idea that a reason's weight varies in proportion to the strength of the desire on which it depends and the extent to which it promotes the object of desire. Instead, claims Schroeder, the weight of a reason depends on reasons to place weight on it. Step two is to argue that most reasons *not* to place weight on moral reasons are going to be reasons of the wrong kind, so they don't count:

If Ryan can't stand Katie, for example, Ryan may have abundant reasons to place less weight on this reason. But those reasons aren't relevant to its *weight*, because they won't be of the right kind. A reason has a certain weight just in case it is *correct* to place that much weight on it. And correctness is determined by reasons of the right kind. (2007, p. 142)

Earlier in the book, we learn what makes a reason one of the right kind:

The right kind of reasons involved in any activity are the ones that the people involved in that activity have, *because* they are engaged in that activity. So, for example, there are correct and incorrect moves to make in chess. The incorrect moves are ruled out, I think, by reasons to follow the rules of the game. Who has those reasons? Anyone who is playing chess. (2007, p. 135)

So, apparently, in order to find out what reasons are of the right kind in order to confer or detract weight, what we have to ask is: in virtue of what activity does Ryan have the reason to help Katie? Now this is a peculiar question. According to Schroeder himself, Ryan has the reason in virtue of (almost) any possible desire. But clearly I don't have all of my desires in virtue of engagement in activities. Or even if I do, then not in virtue of engagement in rule-governed activities like chess. Suppose I have a desire to call my girlfriend. I don't seem to have this desire in virtue of engagement in any particular activity. Suppose I notice it after having worked on an article for some time. How do the 'rules' governing that activity limit the number of reasons to place weight on my reason to call her? It's far from clear.

Perhaps these problems can be avoided by making a distinction: activities are relevant only to reasons (*not*) to place weight on reasons. So even though my reasons to play chess or call my girlfriend do not stem from being engaged in an activity, the *relevance* of a reason to place weight on certain reasons does stem from the kind of activity I'm engaged in. To illustrate: suppose that I am composing a song and that I have reasons both to decrease and increase its tempo. It seems right that there is something odd about making the tempo of my song depend on the state of my hair. That is, arguably, a reason of the wrong kind to place weight on either reason. So reasons to place weight on my respective reasons to decrease and increase the tempo ought to be reasons that are relevant to *musical* aspects of the song. That is presumably how being engaged in the activity of composition restricts the number of reasons for placing weight on reasons internal to that activity.

This is fine as far as it goes. But how is it supposed to help explain why moral reasons are weighty reasons for anyone? Let me quote a bit from Schroeder:

A reason has a certain weight just in case it is correct to place that much weight on it. And correctness is determined by reasons of the right kind. [T]hat means that they must be reasons that everyone who is placing weight on reasons has, in virtue of being someone who is placing weight on reasons. But the activity of placing weight on reasons just is the activity of deciding what to do. So it is simply the activity that every agent is engaged in. So the right kind of reasons with respect to the correctness on placing weight on reasons are precisely the class of *agent-neutral* reasons. It follows that Ryan's idiosyncratic reasons to place less weight on his reason to help Katie are irrelevant, the wrong kind of reason to determine its weight. (2007, p. 142)

Now what is going on here? I think the argument can be represented as follows:

- (1) What reasons (to place weight on reasons) are of the right kind is determined by the activity one is engaged in.
- (2) The activity one is engaged in in deciding how much weight to assign to reasons is deliberation.
- (3) Since deliberation is what *every* agent is engaged in, reasons to place weight on reasons must be agent-neutral (that is, everyone must share them).
- (4) Therefore, reasons that not everybody shares cannot be reasons to place weight on reasons.

I think that premise (3) in this argument is false (and doesn't follow from (1) and (2)). Any decision what to do involves deliberation. So if Schroeder is right, that means that no matter what one deliberates about, only agent-neutral reasons to place weight on reasons are of the right kind. But that is incredible. Suppose I deliberate about whether to have chocolate ice cream or vanilla ice cream. My reason for wanting chocolate ice cream is that I like it. My reason for wanting vanilla ice cream is that I like it. Suppose these are my only first-order reasons. In the context I am in, I have a reason to place more weight on my liking for vanilla ice cream, which consists of the fact that I had a chocolate ice cream yesterday and I like some variation. This is clearly not an agent-neutral reason, but an excellent one to place more weight on my liking for vanilla.

A general problem with Schroeder's argument is that the kind of activity one is engaged in is sensitive to description, and there may be no unique truth about what description is relevant. But different descriptions may have contradictory implications for the question what reasons are of the right kind (to place weight on reasons). For example: when deliberating about my ice cream, I am engaged in deliberation. But I am also engaged in deliberation about my ice cream. These are two different descriptions of the activity I am engaged in. The first is less specific than the last. Everyone who is deciding about anything at all is engaged in the former, but not everyone who is deciding about anything at all is engaged in the latter. Does that mean that it is *both* the case that only unconstrained agent-neutral reasons are of the right kind to place weight on reasons *and* that only reasons that are shared between people who are deciding *about an ice cream* are of the right kind?

Problems abound. In deliberation, one may ask whether to do one thing (buy an ice cream) or another (help Katie). But it is hard to see how the activity of deliberation could by itself furnish one with reasons either to choose the one or the other *without a precise and substantive description of the nature of deliberation* (its objective). For example: if the goal of deliberation is to realise a majority of one's desires, then it is clear how being engaged in deliberation gives one a reason to place weight on reasons to do the one, rather than the other. But Schroeder's description of deliberation as 'the activity of placing weight on reasons' is much too general to make it clear how being engaged in *that* limits the number of reasons that ought to be taken into account.

This problem is aggravated by the fact that we cannot make a neat distinction between first-order reasons to *A* or reasons not to *A* and second-order reasons to place weight on reasons to *A* or reasons not to *A*. Suppose I have a reason to go and see some opera. Is the fact that I also have to finish an article a first-order reason *not* to see some opera, or a second-order reason to place less weight on my reason to see some opera? If almost anything can count as a reason to place weight on reasons, then almost nothing is going to be a reason of the wrong kind to place (less) weight on reasons in virtue of the fact that one is engaged in the activity of placing weight on reasons. But then it is unclear why Ryan's dislike of Katie is a reason of the wrong kind.

So it seems that Schroeder's move of making the weight of a reason depend on reasons to place weight on it cannot guarantee that moral reasons are necessarily weighty for all people. In this section, I have not taken issue with the thought that lies behind the move: the idea that the weight of a reason can indeed be explained in terms of reasons to place weight on it. But that idea is flawed as well.

5. An alternative to proportionalism?

Schroeder proposes a novel conception of the weight of reasons. As we have seen, he rejects *proportionalism*. He also rejects the view that weight is to be thought of in terms of amounts

of numerical value. For if so, argues Schroeder, it ought to make sense to add up Ronnie's reason to go to the party and Wilma's reason to see a doctor. So he proposes a novel account: the weight of a set of reasons to A for an agent is to be thought of as a relation in which that set stands to another set of reasons for that agent not to A . This relation is the relation of being weightier-than (or being less-weighty-than). Unfortunately, Schroeder failed to give a helpful account of what it *is* for a set of reasons to be weightier-than some other set of reasons.

Essential to Schroeder's account of the weight of reasons is the idea that the weight of a reason depends on reasons to place weight on it. This idea stems from the phenomenon of undercutting defeat in the case of epistemic reasons:

[Suppose you see Tom Grabit] come out of the library, pull a book from beneath his shirt, cackle gleefully, and then scurry off. When you see this, you have some reason to believe that Tom has stolen a book from the library. But if Tom has a twin brother Tim, from whom you cannot visually distinguish him, then the case is more complicated. At first it seems that in this modified version of the case, you no longer have a reason to believe that Tom stole a book. But . . . this can be observed to be incorrect by considering a third version of the case, in which Tom and Tim have a third identical sibling, Tam. In that case, your reason for believing that Tom stole a book is even worse than in the second case, so it can't have gone away entirely in the second case. (2007, p. 132)

Schroeder, perhaps quite plausibly, assumes that there is a point where defeaters and defeaters of defeaters (and so on) run out. When that happens, 'we can go back through the chain and determine the weight of your original visual evidence that Tom stole a book' (2007, p. 137). This running out of defeaters is supposed to constitute the essence of Schroeder's account of what it is for a set of reasons to be weightier than another set of reasons. Schroeder gives a recursive account of the weightier-than relation as follows:

Weight Base One way for set of reasons A to be weightier than set of reasons B is for B to be empty, but A non-empty.

Weight Recursion The other way for set of reasons A to be weightier than set of reasons B is for the set of all the (right kind of) reasons to place more weight on A to be weightier than the set of all the (right kind of) reasons to place more weight on B . (2007, p. 138)

Now, a problem is that Weight Base seems fine, but Weight Recursion unhelpful. For what is it for the set of all the (right kind of) reasons to place more weight on A to be weightier than the set of all the (right kind of) reasons to place more weight on B ? Apparently, Schroeder thinks it is just a question of reapplying Weight Base to the sets of reasons to place weight on reasons:

Weight Base simply tries to characterize what it takes for . . . an explanation of the weight of some reason ultimately to get started. Weight Recursion tells us how, once it is started, it continues to proceed. (2007, p. 139)

If that's right, then Schroeder's account of the weight of reasons seems to amount to the following: a set of reasons A to do E is weightier than a set of reasons B not to do E either if B is empty (there are no reasons not to do E) or if there are reasons to place more weight on A than on B . The latter is the case when the set of reasons to place more weight on B is empty (which is not to say that there might not be reasons to place *some*

weight on *B*; it's just that there can't be a set of reasons to place *more* weight on *B* than *A*, if we know that there is a set of reasons to place more weight on *A* than *B*).

But this is still completely unhelpful. For how are we supposed to *know* that the set of reasons to place more weight on *B* is empty? In order for us to know this, we have to know something about the relative strength of the reasons for doing *E* and not doing *E*. And here's the important point: *it doesn't matter that defeaters run out*, since defeaters don't necessarily take all of a reason's force away. Schroeder himself seems to acknowledge this, when he says that Tam makes my original reason for believing that Tom stole the book 'even worse than in the second case'. By a reason being 'worse' than another, he presumably means that it does not have equal weight. But, apparently, my reason for believing that Tom stole the book still had *some* weight (despite Tim's undercutting it). This, however, opens up the possibility that the amount by which a defeater reduces the weight of a reason to *A* is not enough to make the latter weaker than some other reason not to *A*. If so, a reason R_1 to place weight on some reason R_0 may still be stronger than a reason not to place weight on R_0 , despite the presence of a defeater for R_1 .

It is hard to give a simple example of this since reasons come from many corners (the more so, since, as argued at the end of Section 4, we cannot make a neat distinction between first-order reasons to *A* and second-order reasons (not) to place weight on reasons to *A*). For most simple examples with a reason not to place weight on some other reason, it could be argued that there *is* something which functions as a defeater for the first. But even though it may be hard to provide an example, the conceptual point still reveals that the weight of a reason is not something which exclusively consists in the presence or absence of defeaters. And that implies that in understanding what it is for a reason to be weightier than another, it doesn't help to be told that that one is less weighty, for which the set of reasons to place more weight on it is empty. For whether this is so depends on the relative weights of the reasons pro and con. But we cannot presuppose that we know what this consists in.

It seems, then, that Schroeder has not managed to replace proportionalism with a viable alternative.

6. Conclusion

I have never been a fan of error-theories. It strikes me as queer to suppose that moral language is systematically false. Especially if the reason is that it would refer to mysterious entities. It makes morals too much like quasi-scientific discourse. Nevertheless, it is clear that

When we morally condemn a criminal we do not first ascertain the state of his desires. Were we to discover that his desires were well-served by his crimes, perhaps even to the point of wanting punishment, we do not respond 'Oh, well I suppose you ought to have done it after all.' (Joyce, 2001, pp. 42–43)

But if moral statements entail statements about reasons, then it seems those reasons ought to be independent of desires too. Humeanism denies this. It would be a tremendous feat if one could show that Humeanism allows for genuinely agent-neutral reasons. It may allow our judgment about the criminal to be true without denying that it entails a judgment about reasons. Furthermore, it may allow our judgment to be true without committing ourselves to mysterious entities. But, as we have seen, the prospects aren't good. Schroeder's sophisticated theory does not deliver what it promises. It is false that moral reasons can be grounded in (almost) any possible desire. It is also false that even if they could, they would be weighty for anyone.

I think the time is ripe for a reassessment of Humeanism, even by sceptics of agent-neutral reasons. The phenomenon that Joyce describes is hard to square with the view that moral statements entail statements about Humean reasons. For if the latter refer to the desires of *the agent*, why does *the speaker* disregard them in his moral judgments? Why would a judgment that isn't attuned to the desires of others *entail* anything about them?

However, a rejection of Humeanism needn't imply going in for unexplained normative entities (or an error-theory). The phenomenon described by Joyce seems compatible with the view that moral statements implicitly refer to the speaker's (as opposed to agent's) values (perhaps like Dreier, 1990). So perhaps the data have been wrongly interpreted as indicating strongly objective truth conditions (see Evers, 2008; Finlay, 2008). Perhaps speakers' estimates of the moral reasons that are there depend on their commitment to a certain type of values. This would at least explain the intimate connection between moral judgments and judgments about reasons. Speakers don't retract their judgments in light of the discovery that someone lacks certain desires, because their judgments are based on their *own* rather than another's values.

This view may also be able to explain how speakers can truly believe that the criminal has *strong* moral reason to abstain from his criminal behaviour, even though he lacks desires that are furthered by it: the truth conditions of judgments about the strength of a reason may equally depend on *the speaker's* normative perspective (or perhaps even the strength of the speaker's desires concerning what the agent does).

One of the main challenges for such a view is to explain why a difference in desires sometimes seems to make a difference in reasons. For Ronnie, the fact that there will be dancing at the party is a reason to go, but not for Bradley. The question is whether we can explain *why* we make certain reason claims depend on information about the agent's desires. But, on the hypothesis that moral statements (and therefore statements about reasons) implicitly refer to the speaker's values, we may be able to explain this as follows: a speaker judges that Ronnie has a reason to go to the party because he is committed to the value of desire satisfaction (*ceteris paribus*, of course).

Here is an argument in favour of this view: suppose we don't think that the dancing at the party is a consideration that ought to motivate Ronnie. In such a case, we would presumably not say that it is a reason for him to go. Or if we *are* prepared to say this, what we mean seems to be this: it is a consideration that Ronnie *regards* as a reason (or something he is motivated by). So we need to distinguish between normative and motivating reasons. The latter may depend on the agent's desires, but the former do not (or only in the way I have explained).

I may be wrong about the details. But Schroeder's failure to account for agent-neutral reasons contributes to the case against Humeanism. It should inspire even the metaphysically parsimonious to search for an alternative theory of normative practical reasons.

Acknowledgements

I am grateful to Krister Bykvist and two anonymous referees of *Philosophical Explorations* for helpful comments on a previous version of this paper. I would also like to thank the Arts & Humanities Research Council, the Royal Institute of Philosophy, the Nuffic, the Prins Bernhard Cultuurfonds, the VSB fonds, the Stichting Fundatie van de Vrijvrouw van Renswoude te 's Gravenhage, the Stichting Hendrik Muller's Vaderlandsch Fonds, the Stichting tot Financiële Ondersteuning Niet-Gesubsidieerde Projecten voor Mensen in Ontwikkeling and Jesus College, Oxford for their generous support of my studies in Oxford.

Notes

1. Although Schroeder doesn't qualify it, (K) should presumably be read as saying: 'The fact that Katie needs help is a reason to help her for anyone who is in a position to do so'.
2. I am indebted to an anonymous referee of *Philosophical Explorations* for pointing this out.

Notes on contributor

Daan Evers is a DPhil candidate in Philosophy at Jesus College, Oxford. He holds degrees in philosophy from Oxford (BPhil) and Utrecht University (BA and MA), and is currently a (non-stipendiary) lecturer in philosophy at Oriel College, Oxford. He specializes in metaethics.

References

- Dreier, J. 1990. Internalism and speaker relativism. *Ethics* 101: 6–26.
- Evers, D. 2008. *Moral language and objectivity*. University of Oxford, Unpublished manuscript, Retrieved from <http://oxford.academia.edu/DaanEvers>
- Finlay, S. 2008. The error in the error theory. *Australasian Journal of Philosophy* 86: 347–369.
- Hare, R. 1981. *Moral thinking. Its levels, method and point*. Oxford, UK: Clarendon Press.
- Harman, G. 1985. Is there a single true morality? In *Morality, reason and truth*, ed. D. Copp and D. Zimmerman, 27–48. Totowa, NJ: Rowan and Allanheld.
- Joyce, R. 2001. *The myth of morality*. Cambridge, UK: Cambridge University Press.
- Kant, I. 1998. *Groundwork of the metaphysic of morals*, (M. Gregor, Trans.) (Original work published 1785). Cambridge: Cambridge University Press.
- Mackie, J. 1977. *Ethics: Inventing right and wrong*. London: Penguin.
- Parfit, D. 2008. *On what matters*. University of Oxford, Unpublished manuscript.
- Schroeder, M. 2007. *Slaves of the passions*. Oxford, UK: Oxford University Press.