

Being Realistic about Reasons

T. M. Scanlon

Lecture 5: Normative Structure

In my first lecture I listed seven questions about reasons that seemed to require answers. These were:

Relational Character: Reasons are reasons *for* an agent. How is this relational character to be understood?

Ground of Truth: In virtue of what are claims about reasons true, when they are true?

Does the idea that claims about reasons are true or false, independent of our opinions about them, and that truths about reasons are irreducibly normative, have unacceptable metaphysical implications?

Supervenience: How are facts about reasons related to natural facts? They are not entailed by natural facts, but they can vary when natural facts vary and cannot vary when natural facts do not. This may seem puzzling.

Knowledge: If claims about reasons are claims about irreducibly normative truths, how we can come to know truths of this kind?

Practical significance: Judgments about reasons play a different role than other beliefs—such as beliefs about the natural world—in practical reasoning and in the explanation of action. How can they play this role if they are beliefs?

Strength: Reasons have varying strengths. The reason to turn the wheel of the car in order to avoid hitting a pedestrian, for example, is a stronger reason than the reason to go on listening to enjoyable music. What is strength, and how do we determine the strength of different reasons?

Optionality: Some reasons seem to be “optional”—merely reasons it makes sense to act on if one chooses—whereas other reasons are normatively conclusive—reasons it does not make sense *not* to act on. How should this difference be understood?

In the intervening lectures I have offered responses to five of these questions. I will begin this final lecture by reviewing these answers and then say something about the remaining two.

I have taken what may seem to be a short way with the first question, by asserting that the basic element of normative judgments is itself a relation, the relation $R(p, c, a)$, which holds when a consideration p is a reason for an agent in circumstances c to do action a . This thesis seems to me to have considerable explanatory value. The relational character of reasons is most likely to seem puzzling if we focus on reasons *themselves*, that is to say the states of affairs, p that stand in this relation to agents and their actions. If we take the basic normative claims to be apparently non-relational claims that these things “are reasons,” or apparently non-relational claims that certain things “are good,” the question then naturally arises what these normative facts have to do with particular agents. (This puzzlement lies behind Chris Korsgaard’s caricature when she says that according to a realist view we notice reasons “as it were, wafting by.”) The idea that the basic normative element is a relation avoids this puzzlement. It does not seem to me *ad*

hoc, but quite natural once one thinks of it. It also explains other features of normative truths.

In particular, it provides the basis for a plausible interpretation of the relation between “facts and values,” and explains what seems to me to be the more puzzling aspect of the phenomenon of supervenience, namely the fact that many normative truths co-vary with non-normative truths even though they are not entailed by them. More exactly, those normative facts *that vary at all*, co-vary with non-normative facts even though they are not entailed by them. This relation holds in virtue of the truth of what I called pure normative claims, which assign normative significance to non-normative facts. These pure claims themselves, however, do not vary.

I also took what may have seemed to be a short way with the question of motivation, or, as I would call it, the question of normative significance. It is, I said, part of being a rational agent that one’s beliefs about what one has reason to do generally influence one’s subsequent behavior, and can explain that behavior. Building this connection between belief and action into the concept of a rational agent may seem question begging. But it should not seem so. All of the non-cognitivist views that are alternatives to mine explain the relation between normative judgment and action by appealing some psychological ideal type: an agent who (normally) responds to the imperatives he or she issues, an agent who (normally) carries out the plans he or she has made, and so on. The difference lies only in the particular psychological ideal type appealed to. There seems to me no reason to prefer these alternatives to the one I propose, given that, as I argued in Lectures 2 and 4, there are no metaphysical or epistemological objections to taking normative judgments to be capable of truth and possible objects of

belief and knowledge. The fact that a cognitivist account provides a more natural and attractive interpretation of interpersonal argument about reasons is, moreover, a ground for preferring it.

In discussion after my third lecture Ralph Wedgwood and John Broome offered an objection along the following lines. A normative belief that one has conclusive reason to act does not always lead to subsequent action. This indicates that in those cases in which action does ensure the explanation of the action involves the presence of something more, in addition to this belief—an additional element of motivation. I can see the appeal of this argument, but it seems to me to get things backwards. A rational agent is so constituted that his or her normative beliefs generally lead to action in accord with these beliefs. But we are not perfectly rational: things can go wrong. We can be distracted, depressed, too focused on certain pleasures even though we can see that they are not good reasons. *Akratic* behavior is explained in these ways as a malfunction of the apparatus, as a result of which what would normally be sufficient to produce action does not do so. What goes wrong might be described as a lack of “motivation,” but what this means is just a failure of the normal functioning of the agent’s rational capacities, not the absence of some needed element in addition to normally functioning capacities and a normative belief.

I stated the question of the “ground of truth” in the following way: “In virtue of what are claims about reasons true, when they are true? Does the idea that claims about reasons are true or false, independent of our opinions about them, and that truths about reasons are irreducibly normative, have unacceptable metaphysical implications?”

I argued in Lecture 2 that the idea that there are irreducibly normative truths about practical reasons does not have unacceptable metaphysical implications. In my view there are normative facts only in the minimal sense in which these are “merely the reflection of true normative beliefs.” More strongly, my view of the relation between normative and non-normative facts implies that no non-normative facts could be the “ground of truth” for normative truths. Non-normative facts can be the things “in virtue of which” normative claims are true only by being the facts which *are* reasons, not by “grounding” in any sense the fact that these things are reasons. The fact that this metal is sharp is a reason for me not to press my hand against it. But the fact that this is a reason has no non-normative ground, and no normative ground either except possibly further normative facts about my reasons for wanting my hand not to be cut.

This leaves two further questions about the truth of normative claims. The first is what I called in Lecture 4 the question of determinateness: Do claims about practical reasons have determinate truth values? The second is in what sense these truth values are “independent of us.” I have said repeatedly in earlier lectures that the only way we have for discovering normative facts is by carefully considering what seem to us to be normative truths through the kind of process that Rawls called the Method of Reflective Equilibrium. Insofar as this method is just one of seeking coherence in our own attitudes, it may be doubted that it can yield justified belief about matters independent of us. It will be helpful in answering this question and the question of determinateness to say a more about the nature and status of this process.

As Rawls describes the method, it proceeds as follows. One begins by identifying a set of considered judgments about the subject. These are the judgments that seem

clearly to be correct, and seem so under conditions that are conducive to making good judgments of the relevant kind: when one is fully informed about the matter in question, thinking carefully and clearly about it, and not subject to conflicts of interest or other factors that are likely to distort one's judgment. These judgments may be of any level of generality: in the case at hand, judgments about what is or is not a reason in specific cases, for example, or more general claims about what kinds of things can be reasons, or what it is for something to be a reason.

The second stage of the process is to try to formulate principles that would "account for" these judgments. This means principles such that, had one simply been trying to apply them, rather than trying to decide what seemed to one to be the case, one would have been led to this same set of judgments. Since one's first attempt to come up with such principles is unlikely to be successful, there is a third stage, in which one decides how to respond to the divergence between these principles and considered judgments. Should one modify the judgments that the principles fail to account for, or modify the principles, in hopes of achieving a better fit?

It is likely that some accommodation of both of these kinds may be required. One is then to continue in this way, working back and forth between principles and judgments, until one reaches a set of principles and a set of judgments between which there is no conflict, the state of so-called reflective equilibrium. It should be emphasized that this is not a state we are currently in, or likely to reach. It is rather an ideal, which we struggle to attain.¹

¹ In "Reply to Habermas," Rawls says that reflective equilibrium is "a point at infinity we can never reach, though we may get closer to it in the sense that through discussion, our

This method is open to two interpretations. On what might be called the descriptive interpretation its aim is to come up with an accurate description of our views about the subject in question. Rawls sometimes describes it in this way, as a process giving an account of “our moral sense” or “our sense of justice.” So understood, the process would not appear to yield conclusions about anything independent of us.

In Rawls’s earliest statement of the method “our” seems to be understood collectively, rather than individually. The considered judgments with which we begin are not those of a particular individual but the considered judgments of “competent judges.” If the process is understood in this way, then modifying the considered judgments with which it began in order to better fit the principles one has devised to describe them would seem like fudging the data. This is particularly so if the judgments in question are not just those of the person carrying out the process but also those of other “competent judges.”

On the alternative interpretation, which is the one I have in mind, the method of reflective equilibrium is a way of making up one’s mind what to think about a subject. This makes the process, at least initially, a first-person enterprise, and casts the matter of revising one’s considered judgments in a different light.² When one finds that there is a conflict between a considered judgment and a principle that seems to one to have some support, one has to decide what one really believes. This is not simply a matter of weighing the brute intuitive plausibility of the principle against that of the judgment with which it conflicts. Discovering that a considered judgment conflicts with a plausible

ideals, principles, and judgments seem more reasonable to us and we regard them as better founded than they were before.” (PL, 385)

² The considered judgments of others are not *irrelevant*. They are relevant in calling one’s attention to alternatives, raising the question, “Why shouldn’t that be among my considered judgments?” But the fact that others do in fact accept it, and whether or not they could be persuaded to change their minds, does not matter.

principle can tell us something new about both. As we saw in the case of Normative Desire theory, counterexamples can undermine the plausibility of a principle by calling our attention to weaknesses in the thinking that led us to it. The reverse can also happen: when we see what general principles a judgment conflicts with, and see what a general principle supporting it would have to be like, this can undermine the initial plausibility of the judgment itself. As Rawls wrote in a striking paragraph, , “Moral philosophy is Socratic: we may want to change our present considered judgments once their regulative principles are brought to light. And we may want to do this even though these principles are a perfect fit. A knowledge of these principles may suggest further reflections that lead us to revise our judgments.”³

The conclusions of this process are beliefs about some subject matter, “independent of us” that have withstood critical examination. If we have carried out this process as far as we can, or as far as seems necessary to answer some particular question, it might be said that the conclusions we have arrived at are supported by a combination of their intrinsic plausibility and their coherence with our other beliefs. This is true in a way, but also subject to misinterpretation, insofar as it suggests that these two components—plausibility and coherence—are independent. Intrinsic plausibility is not a constant throughout the process. As I have mentioned, many of the considered judgments with which conclude will be different from those with which we began. Some of the latter will have been, if not rejected, then at least reinterpreted and understood in a different way. Whether this is so or not, the “plausibility” of the judgments with which we conclude—our confidence in them—will be in part a function of their having survived this process.

³ *A Theory of Justice* (First Edition.), p. 49 (omitted in Revised edition.)

Although the end result of the process might, ideally, be a set of coherent beliefs (i.e. beliefs that are consistent and mutually supporting) “coherence” is not a good description of what we are seeking in carrying it out. What we are doing is rather testing and evaluating our beliefs in the light of (what else?) those other beliefs that seem most plausible.

The Method of Reflective Equilibrium is a way of deciding what to believe about a subject. It is not itself an account of the truth about that subject. It would be a mistake to say that the truth about practical reasons is given by those judgments about reasons that we would accept in reflective equilibrium. The judgments one would accept at the end of a process of seeking reflective equilibrium will depend on the judgments one has considered as possibilities, the correctness of one’s decisions about which of them to regard initially as considered judgments, and the correctness of the decisions one makes at each stage about how and whether to revise or reject judgments and or principles that account for them or fail to do so. Its dependence on these substantive questions would make such an account of the truth about a subject matter trivial. The method of reflective equilibrium is a way of coming up with an account of a subject matter, it is not itself not such an account.

For the same reason the Method of Reflective Equilibrium is not a constructivist procedure which constitutes the subject at hand. Even if reflective equilibrium is the method we should use to *decide* what sets there are or what reasons we have this procedure itself is not an account of how the domain of sets, or reasons, is constructed because carrying out that process requires decisions about this very question.

Why should we think about a subject in the way that the method of reflective equilibrium prescribes? In particular, why look for general principles that “account for” our considered judgments? Before answering this question I need to address an ambiguity in the understanding of the method. When Rawls first introduced the method of reflective equilibrium he described it as beginning with considered judgments about what is right and wrong in particular cases. In later versions he expanded the starting points to include those judgments of any level of generality in which we have most confidence. This would include, in addition to judgments about particular cases of right, wrong, justice or injustice, general maxims (promises ought to be kept, slavery is always wrong), and judgments about the kinds of considerations that are relevant or irrelevant to questions of rightness and justice. Rawls even says that in order to reach what he calls “wide reflective equilibrium” we should “consider other plausible conceptions [other than the one we have arrived at] and assess their supporting grounds. Taking this process to the limit, one seeks the conception, or plurality of conceptions, that would survive the rational consideration of all feasible conceptions and all reasonable arguments for them.” (CP 289)

Once the process is described in such inclusive terms it may truly seem that there is no alternative to it. But at the same time, and for the same reason, it may seem that there is nothing to it. The “method” is just to think carefully about a question, taking into account and assessing everything that seems potentially relevant to it. This seems to me largely correct. The distinctive content of the method is negative: in its refusal to give privileged status either to particular judgments or to any class of more general truths, axioms, or a priori principles. Everything is up for grabs.

This brings us back to the question I raised before launching off on this digression: why seek more general principles to “account for” our particular judgments (or, more exactly, to systematize those judgments of varying degrees of generality with which we began)? I will mention two possible aims in doing so: large scale and small scale. The large scale aim is to come up with an overall characterization of the subject, such as that provided by a description like the Iterative Conception of set, or by a set of axioms, or by Rawls’s Two Principles of justice. One might seek such an overall account in order to have a basis for deciding particular questions. Alternatively, one might seek such an account in order to describe the subject in way that would provide a basis for answering general theoretical questions about it. One such question is what I have called the question of determinateness.

There are determinate answers to questions about practical reasons if there are (correct) substantive principles characterizing the domain of reasons from which these answers follow. So, for example, if normative desire theory were correct, then questions about reasons for action would have determinate truth values (whether or not we had the information about desires and what promotes them that would be required to figure out what these truth values are.)⁴ As I have said, however, it seems to me that the method of reflective equilibrium does not lead to support of this view. Determinateness could be established in some other “wholesale” manner by establishing some other account of

⁴ My position is thus like Dummett’s anti-realism in allowing for the possibility that some statements about practical reasons may lack truth values because there are no principles determining what these values are. It differs from Dummettian anti-realism in allowing that (as noted in the parenthetical remark in this sentence) if there are such principles then statements following from them have determinate truth values even though we may have no way of finding out what these are. In this it is more like what I called in Lecture 4 a “moderate” view.

reasons of similar generality. I have not ruled this out, although I do not know what such a theory would be. I doubt that there is any such account, because it seems to me that reasons for action are too diverse to be accounted for by any single substantive account.

In the absence of any such general account of practical reasons, determinateness can only be established piecemeal: by determining the truth values of low level generalizations about reasons. Here our thinking generally employs a small scale version of the method of reflective equilibrium. But this method is employed as much to clarify what our reasons are as to verify that they are reasons. Recall the example I used in earlier lectures: the fact that a piece of metal is sharp is a reason for a person not to press his or her hand against it. It is one of our considered judgments that this is so in many circumstances, but it is not always so. It is another considered judgment that someone can have good reason to try to cut his hand. So the conditions *c* in our original statement have to be stated so as to preclude such cases. Searching for the relevant *c* is a matter of trying to find a (low level) principle that “accounts for” our considered judgments about when we have reason not to press our hand against the metal. But this could just as well be described as a process of trying to figure out what these implicitly general considered judgments really are.

Most of our day to day thinking about reasons for action takes place within a framework of accepted judgments of this kind. When it seems to us that, as David Wiggins might put it, “There is nothing to think but that *p* is a reason,” we may be saying (as perhaps we are in the case of pain) that the judgment that *p* is a reason seems to us, on reflection, so obvious that we cannot imagine its being mistaken, or we may be

recognizing that p's being a reason clearly follows from more general precepts about reasons that have this status for us.

These more general precepts are commonly put as views about what is good, right and wrong, virtuous or vicious or in terms of other "thick ethical concepts." Even if I am correct in believing that the normative content of these notions is to be understood, ultimately, in terms of the basic normative idea of a reason, focusing on isolated claims of the form $R(p, c, a)$ may make practical thinking appear less determinate than it really is by obscuring the normative context that these concepts supply. My focus on such judgments is still appropriate, because it helps to clarify the relation between normative and non-normative facts. More to the point, judgments made with these thick concepts seem to have determinate truth values because we are holding fixed the judgments about reasons that they embody. They have empirical content, as Tim Williamson pointed out in discussion last week. But they have this content because certain empirical facts are assigned normative significance by normative judgments that are presupposed when we use those concepts. These judgments are, however, open to the same generally epistemological worries as the particular judgments about reasons that I have been focusing on. To put the same point another way: our confidence in these judgments depends, ultimately, on the same kind of reflective equilibrium process that I am now discussing.

Are facts about practical reasons, as I have described them, independent of us? "Independence" can mean a number of different things, which need to be distinguished. Facts about reasons are relational facts, about the significance of certain natural facts for

rational agents. They thus are not facts “out there” having nothing to do with us. But this is not the kind of dependence that might seem problematic.

It is sometimes said that facts about reasons are dependent on us because there would be no reasons in a world without rational agents. Certainly there would be no point to talking about reasons in the absence of rational agents for whom these are reasons. But it could still be true of certain facts that if there were agents (in the relevant circumstances) then these facts would be reasons for them, if they existed. If we can say, sensibly, that it would be a bad thing for the world to be a certain way even if there were no rational agents in it, this is just a claim about how *we* have reason to feel about such a world, or what we have reason to do insofar as we might affect whether there would be such a world or not.

The kind of independence of us that normative facts are supposed to have, and do have on my view, comes to at least this: they are facts that we, individually, could be mistaken about. Could we all, collectively, be mistaken about certain reasons? We might all fail to notice some relevant feature of a situation, and all die before learning any better. If this could happen through descriptive error or lack of information, why not through normative error as well?

Reasons do depend on us in one further way, however. If our reflective judgments about reasons showed no tendency to converge over time on a stable set of opinions (if our “corrections” were likely to be reversed the next time we thought about the matter, and to oscillate back and forth) then thinking about reasons would be pointless. Similarly, if different agents did not have any tendency to converge in their well-informed reflective judgments about practical reasons, then this would make the practice of offering each

other reasons and arguing about what reasons we have pointless, and would cast doubt on our assumption that there are truths about reasons. This does not mean, however, that our thinking certain things are reasons *makes* them so. Convergence of opinion on a normative claim does not make it true (although it can be a reason to think it true.)

Let me turn now to the question of strength. I have so far been discussing the normative relation $R(p, c, a)$, understood as the minimal claim that p is *a* reason for an agent in circumstances c to do a . This leaves entirely open whether it is a conclusive reason, that should settle the matter of whether to do a , or merely a consideration that it makes sense to take into account in considering whether to do a . Only the latter is claimed, although the former is left open. So understood, R is the relation we use in classifying various factors as those that should play a role in deciding whether to do a as opposed to those that are irrelevant.

In order to express the idea of one consideration's being a stronger reason than another, we need a normative relation that takes account of the normative interaction between different considerations. The relation $R(p, c, a)$, as I have interpreted it, already expresses some interaction of this kind, in virtue of the middle term in this relation. If, for example, q is a factor that undercuts p as a reason, this will be reflected in the fact that $R(p, c, a)$ can be true but $R(p, c', a)$ false, when c' includes q but c does not.

But *outweighing*, the relation that occurs between one consideration, which is a reason to do a , and another consideration that is stronger reason not to do a , is a different matter from undercutting. Even if q is a stronger reason not to do a than p is a reason to do a , and c includes q , p remains a reason to a , just not a good enough reason. So let $R^+(p, c, a)$ be the relation "p is a sufficient reason for an agent in c to do a ." Using this we

can express the relation of strength. If p is a reason to do a and q is a reason to do not a , then q 's being a stronger reason than p just comes to the following: $R(p, c[p,q], a)$, $R^+(q, c[p,q], \text{not-}a)$, and $\text{not-}R^+(p, c[p,q], a)$, where ' $c[p,q]$ ' indicates that c includes both p and q .

Using the same relation, we can also express the idea of an "optional" reason. A reason p is an optional reason to do a if there are considerations q such that $R^+(p, c[p,q], a)$ and $R^+(q, c[p,q], \text{not-}a)$. Note that cases in which reasons are optional are not cases of indeterminacy as I have been discussing it. They are not cases in which claims about reasons lack definite truth values, but rather a particular configuration of such values.

I have shown how to *express* the ideas of relative strength and optionality. There remains the question of the kind of thinking through which we arrive at conclusions of this form.

Desire theories and hedonistic theories offer extremely general accounts of reasons and their strength. They also suggest a methodology for assessing reasons and their strengths that is atomistic: to determine the strength of a reason, on either account, we focus on *it*: on the desire or the amount of pleasure that supports it. If, after rejecting these theories, we retain this methodological picture we may seem to be maintaining that reasons have an ineffable property of strength, which we can assess just by focusing carefully on *them* in the right way. This kind of case-by-case intuition seems implausible.

This problem arises from hypostasizing reasons in an implausible way. Reasons are not a special kind of abstract object. The basic normative element is the relation R (and its relatives such as R^+), and claims made using this relation all have a kind of generality: that p is a reason for a for any agent in circumstances c . Judgments about

relative strength (“intuitions” of relative strength, if you want to put it that way) are judgments about the truth of general claims of this kind, specifically of the kind I outlined above, using R and R^+ .

In this respect my view resembles Kant’s. Accepting that $R^+(p, c, a)$ is what Kant called adopting a maxim, the policy of taking p to be sufficient reason for doing a in circumstances c . Kant seems to me to be correct in holding that maxims (at least in the normative form in which I am interpreting them) are the fundamental elements of practical reasoning. (Whether maxims are also the basic objects of moral appraisal is another question.) As I am interpreting them, maxims express judgments about the relative strength of reasons. Perhaps Kant understood them differently, merely as policies, or laws given to oneself.

Where I clearly differ from Kant, however, is on the question of how we should decide which maxims to adopt, or, as I would put it, which instances of $R^+(p, c, a)$ are correct. According to Kant, as I read him, a person (correctly) decides which maxims to adopt on the basis of his or her inclinations, under the constraints of the Categorical Imperative. I do not accept Kant’s arguments for the Categorical Imperative as a fundamental principle of practical reasoning, and I do not want to say that the strength of reasons is in general determined by the strength of the agent’s desires. Even proponents of desire theories find this implausible.⁵ So something more needs to be said. I do not have a general answer to this question, but I will mention a few patterns of thought leading to conclusions about strength.

Consider first what I will call Dominance cases. If continued life will offer pleasures of intellectual companionship reading for many days, as well as other things

⁵ Cite Schroeder

worth wanting, then it does not make sense to give up life for the sake of one day of these pleasures or reading and intellectual companionship.. That is to say, the fact that doing a will preserve my life is a stronger reason than the fact that doing a' would bring me these particular pleasures of duration x, but no life beyond that. This "tie breaking" argument presupposes that the reasons for the two incompatible actions "balance out," so that the additional reason settles the matter. This already involves a comparative element. I have tried to minimize this element by taking cases in which the initial "balanced" reason providing considerations are exactly the same, but it remains nonetheless.

Quantitative cases are a special type of Dominance cases. In such cases, what gives one a reason to a is the presence of some feature that comes in amounts. So one has a stronger reason to do what will yield more of this feature than one has to do what will yield less. That is to say, one has an additional reason to do this, provided by the additional amount of the crucial feature that it will yield. Quantitative cases come closest to fitting with the "upward" explanation offered by hedonistic theories. Indeed, one reason sometimes given for finding hedonism appealing is the idea that it would provide just this basis for rational decision. The interesting question is how to explain cases of relative strength that do not fit the dominance model.

One kind of such reason is strictly instrumental: that is, it has to do with the causal consequences of following a policy of a certain kind. If one has reason to achieve an end that can only be achieved by following a certain policy, and if one will not follow that policy if one gives considerations of type a priority over considerations of type b, then one has reason to adopt a policy that rules this out, by holding that considerations of type a are outweighed by those of type b in situations of the relevant kind. Policies

regarding exercise and dieting are familiar examples. The case for adopting such policies already contains a comparative element however. One has sufficient reason to adopt a policy of skipping dessert in order to lose weight only if one has reason to prefer life in which one is thinner but has missed out on many sweets to a life in which one is heavier but has had these enjoyments. This is naturally seen as reflecting the fact that one has stronger reason to be thinner than to have these enjoyments.

I will mention one other class of cases, just to distinguish it from these. Particular relationships with others may require holding certain judgments of strength. Being a friend requires thinking that the friend's interests give rise to stronger reasons than one's own convenience, and in most cases stronger reasons than the comparable interests of strangers. Being related to others by common commitment to a cause or value, requires taking that consideration to give stronger reason than other things do.

The point of these examples is just to indicate how our views about the relative strength reflect a larger set of normative attitudes rather than being isolated reactions to relative quantities of some normative, or non-normative property.

I have suggested that we assess our beliefs about reasons for action and decide whether they are correct by engaging in a process of seeking reflective equilibrium among our normative beliefs. But different people, applying this method conscientiously, can reach different conclusions. In cases in which there is such disagreement, how can the kind of correctness that this process can deliver really matter? Isn't some further kind of correctness required? To answer this question it will help to consider a particular matter about which people disagree.

Does the fact that a certain form of behaviour is required by the traditions of my society give me a reason to engage in it? Suppose I think at the outset that it does not. Traditions can be bad. When they are not bad, I think, the reasons we have to do what they require depend simply on their merits.

How do I decide whether this is the case, or whether there is some further tradition-based reason that I am missing? I consider various examples in which it seems that this might be the case, and consider what the reason in question might be and how it might be explained. The fact that a certain form of behaviour is traditional in a society may mean that many, perhaps most people in that society would be disapproving, perhaps shocked at my failure to follow it. If I have reason not to shock, then this is a reason (whether the society is mine or not.) More positively, if I have reason to want to be understood by most people in a society as expressing a certain attitude, such as love and respect for a recently deceased person, and sadness that she has died, then I will have reason to do whatever is the traditional expression of those attitudes.

But the suggestion I began with, that the fact that something is traditional in my society can be a reason to do it, seems to aim at something more fundamental than these instrumental and communicative reasons. What might that be? Here is another thought. There is a special pleasure in doing things with other people with a shared sense of the meaning of these activities for all of you. This is not only pleasant, but a pleasure worth seeking. Moreover, in some cases the shared sense of meaning involves the idea that these activities have been and will continue to be repeated over time. Part of the significance, and value, lies in this repetition and intention to repeat as an expression of enduring connection among the varying participants over time. Seeing the existence of

this practice as giving one no reason to continue it is incompatible with sharing in the good involved—incompatible with being related to the other participants in the way that this good requires. This does not mean that one necessarily has reason to continue the practice. That depends on whether one has reason to continue in the relation it involves. Some shared activities are bad even if the participants enjoy and are ‘bonded’ by them. This fits with the thought I mentioned at the outset, that not all traditions give reasons.

I have gone through these steps to illustrate the way in which I have said we should assess reason claims. It also illustrates what I said earlier about ‘strength.’ How strong is the reason to continue a practice, understood in the way I have just proposed? This is determined at two levels. The first has to do with the understanding of the parties: in order to have the significance they attribute to the practice, what considerations does the reason to continue the practiced have to override? If it can be overridden by the slightest inconvenience, then it does not seem very meaningful. Or, to put the point in reverse, it may be that if I take my reason to continue the practice to be so slight as to be overridden by the slightest inconvenience, then I cannot be seen as a member of the group, sharing in its good in any meaningful sense. A meaningful group is likely to involve more serious commitment. So the first answer to the question of strength is internal to the practice in question. But this answer to the question of strength is dependent on the assumption that one has reason to be a member of the group to begin with. Groups can involve practices that are tyrannical in their demands. In which case one may have good reason not to consider oneself a member.

The same point can be put in terms of optionality. Is the reason one has to continue a practice an optional one (something there is reason to do, but just one reason

among others—one might go to a film instead)? Sharing in the good of participating in that practice is likely to preclude seeing the reason to continue as optional in this way. But the further question is whether seeing oneself as such a member is itself optional. I believe that with respect to almost all groups (smaller than the moral community itself) the answer is yes. Being in a position to be part of a group of the kind I have described given on the opportunity to have a certain good, but it is one good among many and the reasons for pursuing it, though good reasons, are generally optional in the sense I described. Groups that take the opposite view are tyrannical.

What I have just said seems right to me, but I imagine that much of it is controversial. Indeed, I *hope* that at least some of it is controversial, because I have said it with two purposes in mind. The first is to illustrate the way in which, on my account, we do and should go about answering questions about reasons. The second is to raise a question about this method, namely whether it can deliver conclusions that have the significance that claims about reasons need to have. It may seem that it cannot. What I have described is just a way of determining how things seem to me as far as reasons are concerned. But what is that to anyone else? If others disagree, what use are my conclusions if all they reflect is my own careful thinking? Things would be different, one might say, if my conclusions were grounded in something that they also accept, or must accept. But I have argued that this is not so. Isn't my account of reasons rather solipsistic?

So suppose someone disagrees with me about whether and how the traditions of a society provide the members of that society with reasons to continue them. Or suppose they disagree with me about the strength of those reasons (about when a group becomes

tyrannical in the sense that is a strong reason not to be a member.) How does it matter if my view about this is 'correct'?

First, as I said in describing the method of reflective equilibrium, the fact that someone else has reached a conclusion different from mine can be something that I should take into account in assessing my own view. I need to consider whether this person may have thought of alternatives (e.g., ways in which tradition might be significant) that I overlooked. The fact that she reached a different conclusion may indicate some myopia on my part. But suppose that I have taken this into account (and that she has done likewise) but that the disagreement persists. Does this undermine the significance of my conclusion? The question is: its significance for what?

The first possibility is that it might undermine the significance of my conclusion as grounds for *my own* decision about what I should do. I do not see how the fact of disagreement could have this effect except in the way just mentioned, by calling my attention to some mistake in my reasoning. If after reflection it still seems to me that I have reason to continue a tradition (or that I do not, or that this reason is 'optional') that settles the matter for me. What more could be required to give my conclusion significance as a guide to my own action?

The second possibility is that the fact of disagreement might undermine my conclusion as a basis for deciding what *this other person* has reason to do. Why should my conclusion be authoritative for her? But what I have reason to do in circumstances *c* and what she has reason to do *in those same circumstances* are not independent questions. An answer to one is at the same time an answer to the other. The two of us have different reasons only if our circumstances differ in some relevant respect. In this

case that means only if my *belief* that I have a reason is a condition of my having that reason. If it does not, then unless the fact that she disagrees with my conclusion is evidence for modifying that conclusion, it does not make a difference to what either of us has reason to do (although it will, by definition, make a difference to the answer she gives to that question.)

Third, moving into the area of morality, does the fact that she disagrees with me make a difference to her moral claims about what it is permissible for her to do? In my view the answer to this moral question depends on what reasons a person in her situation would have for objecting to a principles that would forbid what she proposes to do. Suppose, for example that she objects to a principle that would forbid certain conduct by claiming that that conduct is required by traditions of her society and that she therefore has good reason to want to be able to engage in that conduct. I believe, on the other hand, based on my analysis of such reasons, that they are not strong enough to support this claim: they do not make it reasonable to reject the principle in question. The moral question turns on the reasonableness of her proposed rejection, that is to say on the strength of the relevant reasons. I can't answer it except by forming a judgment about this question, and having arrived at such a judgment I have no further decision about how to answer it. So the fact of disagreement is again irrelevant, and the only kind of 'authority' that judgment of this sort needs to have is simply being a correct assessment of the relevant reasons.

This result may be made clearer by making explicit the possibility that third parties may be involved. The reason why it is not reasonable to reject the principle for the reason given, the value of continuing tradition, is that the cost of continuing it to others is

too high. The agent in question and these third parties may disagree about this, because they disagree say, about the value of continuing tradition. To form an opinion about the permissibility of the conduct in question one must take a position on this disagreement.

Things may seem to be different if what is in question is not the permissibility of what some other person does but the permissibility of my interfering with what she does. Why should my conclusion that what she takes to be a reason to engage in her conduct is not really a reason license me in interfering with it? The answer is that it probably doesn't. Even if she has no good reason to engage in this conduct (which does not, I am assuming, harm anyone else) she does have good reason to want to live by her own lights, which gives her reason to reject principles what would permit others to interfere (when no one else is harmed.) The case is thus different from the one previously discussed, but not in a fundamental way, because the impermissibility of interference depends on the conclusion just stated about reasons for wanting not to be interfered with. That is to say, it depends on the correctness of this claim, not on whether people agree about it.

Questions about what we have reason to do are questions we are capable of thinking about in the ways I have described. If one has thought carefully about such a question in this way, and sees no objection to it, then there is nothing further one can do to decide whether it is correct. But even if this is all we can do, it would be complacent to infer from this that it is *enough*. As I have stressed, it is not enough to ensure that the conclusion we have reached is correct. We might have made a mistake in our reasoning. But the question whether we have done this is just another first-order normative question, calling for, and answerable by more reasoning of the same kind. The charge of

complacency that I am imagining claims to go farther: it is that this kind of reasoning might be inadequate because it fails to be in touch at any point with the normative facts.

The idea is that there might be facts about which things are reasons for action that we need to “be in contact with” or “be sensitive to” in some way other than through the process I have described in order for our conclusions about reasons to be sound. But I do not see how there could be facts of this kind that were normatively relevant. Facts about reasons are the sort of thing we can think about and assess through a process of the sort I have described. There is no way of “being in touch with them” except through this very process.

This is not to justify complacently thinking that we have gotten it right in regard to reasons. But the thought that we might not have is a thought internal to this reflective process, and the grounds for concern on this ground are also internal to it: they lie in the weakness and indeterminacy of the forms of reasoning that it involves.

Christine Korsgaard has written that on a substantive realist account of reasons we have nothing more to go on than our confidence that, after thinking about what reasons we have, we have gotten it right. She intends this as a criticism of a view of the kind I have been defending. But I there is something right about what she is saying. We do have nothing to rely on except our best *judgments* about which things are reasons. Often our confidence in these judgments justified; sometimes it is not. But that’s life. There is nothing more that we could have. To be realistic about reasons we must accept this fact.