

The John Locke Lectures 2009

Being Realistic about Reasons

T. M. Scanlon

Lecture 3: Motivation and the Appeal of Expressivism

The cognitivist view I have been defending has two important features in common with Gibbard's expressivism and Blackburn's quasi-realism. It does not presuppose normative properties in the (natural) world. And it interprets claims about reasons as *reactions* to the (natural) world: more specifically claims about the appropriateness of certain reactions to it. Where my view differs most obviously from Blackburn's and Gibbard's is in maintaining that normative claims can be true or false (although more needs to be said about how much this comes to.) My view also differs from these non-cognitivist views in the explanation it offers of the "motivational" power of normative claims. Today's lecture will be devoted to examining these differences: to considering how much they come to and how my cognitivist view can be defended. I will begin with the second difference.

It is often said that an adequate account of reasons must explain how they motivate people to act, but it is not clear exactly what kind of explanation is thought to be required. The term 'motivate' has a dual character. This is brought out by considering the contrast between my view and desire theories, which are generally seen as providing a more satisfactory explanation of the kind in question. On the one hand, when it is said that only desires can motivate, and that beliefs are "motivationally inert" the motivation in question may seem to be a kind of causal

efficacy. But ‘motivate’ also has a rational, or even normative aspect: desires are not only supposed to cause actions but also to “rationalize” them, as Davidson famously said.¹ I take this to mean, at least, that a desire makes an action (believed to promote its satisfaction) understandable, or perhaps even makes it “rational.”

If to rationalize an action is to make it understandable, and even rational, for the agent to so act, then it would seem that an agent’s belief that she had a reason to perform an action, even if it is a *belief*, could rationalize her action just as well as a desire could. So the supposed unique motivational efficacy of desires, and corresponding deficiency of cognitivist accounts of reasons, may lie in the former, causal, idea of motivation.

Desires come to us “unbidden” and we may feel that they “impel” us to action. But this does not mean that an explanation of action in terms of the agent’s desires is a causal explanation in a way in which an explanation in terms of the agent’s beliefs about reasons is not. Beliefs about reasons are also not subject to our will or choice, and can have a demanding quality. The deeper point, however, is that neither in the case of desires nor in that of beliefs about reasons is this experience of “impulse” a direct experience of a cause. This feeling is simply an element of our momentary experience. If such experiences are generally followed by action this is because of some underlying neural mechanism that is equally causal in the two cases and in neither case an object of experience.

Even if the belief that one has a reason to act in a certain way can rationalize that action, that is to say, make sense of it, more needs to be said about how, and in

¹ In “Actions, Reasons, and Causes,” in *Essays on Actions and Events* (Oxford: Clarendon Press, 1980), p. 4.

what sense, such a belief can *explain* action. On my view, this explanation relies on the idea of a rational agent. A rational agent is, first, one that is capable of asking questions about the reasons he or she has for performing certain actions or for holding various other attitudes. Second, a being is a rational agent only if the judgments that it makes about reasons make a difference to the actions and attitudes that it proceeds to have. A perfectly rational agent would always have attitudes and perform the actions that are appropriate according to the judgments about reasons that he or she accepts. A rational agent will generally intend to do those actions that he or she judges him or herself to have compelling reason to do, and believe a proposition if he or she takes him or herself to have good evidence for its truth. A perfectly rational agent will always do these things.

None of us is perfectly rational, but it is appropriate to call us rational agents just in case we come sufficiently close to meeting this standard. When a rational agent does something that he or she judges him or herself to have reason to do, this judgment makes sense of the action in normative terms and explains it: it is what one would expect of a rational agent who accepted that judgment. Presumably there is also a causal explanation of this connection, and of the more general uniformities that I have referred to. But this is another story, for neurobiologists to fill in.

Now compare this explanation with that offered by expressivist theories. The accounts that they offer are like mine in emphasizing the “rationalizing” aspect of motivation rather than the causal aspect. Looking back to earlier non-cognitivists, R. M. Hare wrote that moral judgments must be understood as expressing the acceptance of imperatives because this is the only kind of judgment that is logically linked with

action. That is, the only kind of judgment such that if a subsequent action of an appropriate kind is not performed it follows that the agent spoke insincerely, or did not understand what he or she was saying.² More recently, Allan Gibbard in *Thinking How to Live* analyzes judgments about reasons (not necessarily moral judgments but judgments about what one has most reason to do) as decisions about what to do, or the adoption of plans.

Each of these accounts is like mine in explaining the connection between normative judgments and action by appealing to an ideal of a certain kind of being. What Hare mentions is linguistic competence (understanding the meaning of words) and “logic,” but he seems also to be appealing to a kind of practical consistency (acting in accord with the imperatives one accepts.) In Gibbard’s case the appeal is clearly to an ideal of practical rationality: the acceptance of a normative judgment not only makes sense of (rationalizes) an action but also explains it, because rational agents are so constituted that they generally do what they have decided to do, and generally carry out their plans. If they do not do these things then they are irrational.

Note that all three of these explanations of the connection between normative judgment and action—Hare’s, Gibbard’s, and my own—are “naturalistic.” They all refer to ideal psychological types, but the identification of an individual as an instance of one of these types is an empirical question, and the states of mind that these types refer to are entirely naturalistic psychological states.³ The mental state of judging something to be a reason *refers to* something non-naturalistic—a normative truth—

² *The Language of Morals*, pp. 20, 171-172.

³ See Blackburn, *Ruling Passions*, p. 57-58. Blackburn emphasizes (p. 49) that on his own account valuing is a “natural and naturalistically describable state.” This seems equally true of the normative attitudes involved in the other accounts I am considering.

and this may be the contrast that Gibbard and Blackburn have in mind. But the state itself is just as “naturalistic” a state as adopting a plan or, for that matter, having a desire. But there are important differences between these accounts.

One difference between the non-cognitivist views and my cognitivist account lies in the nature of the ideal type referred to: in my case it is an ideal of rationality that is understood explicitly in terms of the acceptance of judgments about *reasons*. Their accounts avoid this, appealing instead to plans, imperatives, or pro-attitudes. Another difference is that my account construes the acceptance of such judgments as kind of belief—as the kind of thing that can be true or false, correct or incorrect—albeit a belief of a special kind, one that is linked by rational requirements to attitudes other than other beliefs. So we need to consider how these differences matter, and whether they provide reasons for preferring one kind of account over the other.

First, it might be claimed, as Gibbard does claim, that non-cognitivist accounts give a deeper explanation of our normative attitudes.⁴ They explain what it is to judge something to be a reason rather than taking this idea for granted, as I do. It is true that non-cognitivists identify normative judgments with states (pro-attitudes, acceptance of imperatives, plans or decisions about the thing to do) that do not explicitly involve the idea of a reason, but which nonetheless have implications for subsequent action and attitudes of the kind that normative attitudes are supposed to have. The question is whether this identification provides an informative explanation of normative judgments, or changes the subject, by identifying normative judgments with something quite different. This is what Gibbard calls the question of *internal*

⁴ See *Thinking How to Live*, pp. 184-191.

adequacy: whether the expressivist analysis “accounts for everything internal to normative thinking, or everything internal that is intelligible.”⁵

What does internal adequacy require? It requires that an account of normative judgments explain their distinctive significance for subsequent action and other attitudes. This is what I have so far been discussing. In addition, internal adequacy seems to require two related things. The first is that an account of normative judgments should do justice to the thought that these judgments can be mistaken and, if correct, would be correct even if one did not make them. To take an example that Gibbard cites, an internally adequate account of normative judgments should be able to make sense of a person who thinks that it is wrong to kick dogs for fun thinks that this would be so even if he, misguidedly, believed that dog kicking was perfectly acceptable behavior.⁶ Second, an internally adequate account of normative judgments should give a satisfactory account of their use in interpersonal discourse, such as in giving advice, or discussing the justifiability of what someone has done.

These requirements pose problems for expressivist views because those views interpret making a normative judgment as *doing* something (such as expressing an emotion, adopting a plan, or accepting an imperative) that, it would seem, one can only do for oneself, and as doing something that has its normative force and efficacy, even for oneself, only when one is doing it.

Gibbard and Blackburn are aware of these problems and have responses to them. I will focus here on Gibbard’s response. Holding that there is a fact of the

⁵ *Thinking How to Live*, p. 186.

⁶ Gibbard, *op cit*. Blackburn makes a similar point in “Errors and the Phenomenology of Value,” p. 153.

matter, independent of us, about what we ought to do—that, for example, we ought not to kick dogs for fun—is, on Gibbard’s view, a matter of “planning to avoid kicking dogs for fun, planning this even for the contingency of being someone who approves of such fun, and who is surrounded by people who approve.”⁷ But is this intelligible? What does it mean to plan not to do x even for the contingency in which one approves of doing x? It does not seem that one can plan to do something under certain circumstances while knowing or believing that if one were in those circumstances one would not do it.

Similarly, when one judges that someone else has conclusive reason to do something, as when one is giving them advice, one can’t be deciding that they will do it, or planning, for them, to do it.⁸ On Gibbard’s view, what one is doing in such a case is planning to do this thing for the contingency in which one is in that person’s situation in all relevant respects. But it seems strained to speak of planning to do something in a situation that one knows one will never be in.

Gibbard recognizes the strain raised by cases of these two kinds. “Fully fledged planning,” he says, involves both coming to an answer about what to do in certain circumstances and expecting “that that thing is what [one] really will do if the contingency arises.”⁹ Cases of contingency planning of the kinds just mentioned, he recognizes, lack this second aspect. Nonetheless, he thinks that this idea of planning, including contingency planning for situations that will not, or cannot, arise, provides the best interpretation of normative judgment.

⁷ *Thinking How to Live*, p. 186.

⁸ John Hawthorne also raises this difficulty in “Practical Realism?” *Philosophy & Phenomenological Research* 64 (2002), p. 171.

⁹ *Thinking How to Live*, p. 49.

The strain that arises for an expressivist view in cases of these two kinds is related to the “Frege-Geach” problem of interpreting sentences in which normative judgments are embedded in more complex sentences. All these problems arise for the same reason. A central originating idea of expressivism is that mere beliefs could not have the practical significance that normative judgments are agreed to have.¹⁰ To explain the link between normative judgment and subsequent attitudes and actions, expressivists say, these judgments must be understood to involve some more “active” element, such as the adoption of an imperative, a plan, or some other pro-attitude. The problems we are now addressing all arise from the fact that normative judgments occur in contexts in which the person who utters them is not doing any of these “active” things.

These are not special or isolated cases. It is natural to think of judgments about reasons for action as arising in response to the question that an agent asks him or herself in deciding what to do. But judgments about reasons occur just as centrally in interpersonal discourse, in cases in which one is offering advice, or discussing what some third party should do, or, importantly, offering a justification for one’s action.

Consider first the case of advice. When I give someone advice, by telling them that they have good reason to do A, what I express to them is not a decision or plan to do A. Rather, I am urging *them* to make such a decision, or to adopt such a plan, by calling their attention to what I take to be good reasons for doing so. Even more obviously, when I attempt to justify my decision to do A to someone who

¹⁰ As Blackburn writes, “Beliefs do not normally explain actions: it takes in addition a desire or concern, a caring for whatever the belief describes.” *Ruling Passions*, p. 90.

wanted me not to do A, I am not merely expressing my decision, or plan. Rather I am asserting that I have good reason to have made that decision, or plan, and trying to get them to accept that this is so. So the operative normative element in what I am doing is not the element linking acceptance of a normative judgment to subsequent action (the link that non-cognitivism was originally designed to explain.) It lies, rather, at a higher level: it concerns the *adoption* of such a judgment.

The view I am recommending captures this difference. I said at the beginning of this lecture that my view was like expressivist views in understanding normative judgments as concerned with our responses to natural facts rather than with those facts themselves. But I then added that normative judgments, on my view, are claims about the *correctness* or *appropriateness* of such reactions. This difference is crucial.

The idea, $R(p, c, a)$, that it is correct for a person in c to treat p as a reason for doing a provides a plausible interpretation of what a person says when she advises a person in c to do a because p . It is something people can agree disagree about and argue about, whether or not they take themselves to be in c , and something they can make assertions about the implications of using embeddings such as: “If $R(p, c, a)$, then ...” They can do all of these things without making any decisions or adopting any plans, although, if they accept that $R(p, c, a)$ and believe themselves to be in c , then (insofar as they are rational) they will do intend to do a , or at least count p in favor of doing a , and if they do a this will be explained by their acceptance of this judgment.

Of course, one of the main points at issue in this debate is whether such claims of correctness make sense. I will have more to say about this later. But first I want to

consider the interpretation that expressivists offer of what I just called the “higher level” of normative discourse and disagreement about whether one should accept a normative judgment about reasons for action.

What expressivists can say is that the correctness of claims about reasons for action is to be understood with reference to norms governing the formation (or abandonment) of our first-order attitudes. Gibbard, for example, holds that to think that *a* is what someone in situation *c* has most reason to do is to plan to do *a* if one is in *c*. To think that *p* is a reason for a person in situation *c* to do *a* is, on this view, to plan to count *p* in favor of doing *a* for the contingency in which one is in situation *c*. To think that one of these attitudes is *correct*, one might say, is just to accept certain norms for the formation of such attitudes and to believe that those norms support forming, or continuing to hold, those attitudes in the epistemic situation in which one now finds oneself.¹¹

Applying this to interpersonal discussion about reasons, we can say that to recommend that another person who is in situation *c* should do *a*, or that such a person should count *p* as a reason for doing *a*, is not only to plan (for the contingency in which one is in *c*) to do *a*, or to count *p* in favor of doing *a*. It is also to express one’s acceptance of norms governing such attitudes and one’s belief that these norms support having the plans just mentioned for behavior in *c*, if one is in the agent’s current epistemic situation.

Similarly, someone who offers *p* as a justification for his doing *a* in situation *c* not only expresses his own acceptance of a plan to count *p* in favor of *a*, but also,

¹¹ Blackburn, “*Ruling Passions*,” p. 293, Gibbard, *Wise Choices, Apt Feelings*, Chapters 8,9, esp. pp. 168ff.

according to this view, expresses his acceptance of higher order norms governing such attitudes, and his belief that these norms support holding an attitude of counting p in favor of a in situation c for someone in the epistemic situation of the person who is being urged to accept this justification.

This analysis provides a more satisfactory interpretation of interpersonal discourse about reasons than the simpler expressivist account I considered earlier, which relied only on lower order attitudes, such as plans about what to do in certain situations or about what considerations to count as weighing in favor of these actions. But this more complex account has limitations similar to those of the simpler version.

These limitations are brought out when we consider the possibility that the person to whom you offer a justification for your action may reject the higher order norms to which your justification appeals. In offering that justification, according to the expressivist account, you are expressing your acceptance of these higher order norms. The other person claims not to accept them. But, intuitively, in addition to expressing your acceptance of the norms you are appealing to, you are claiming that these norms themselves are correct. This claim can also be given the same expressivist interpretation. The expressivist can say that in claiming that these norms are correct you are just expressing your acceptance of yet higher order norms which, you believe, support holding these norms in the epistemic situation in which you and your interlocutor are placed. Since there can be disagreement about these norms in turn, the threat of regress is clear. This regress can be avoided only by claiming, at some level, that the relevant norms are *correct* in a sense that is not to be cashed out in terms of your acceptance of yet higher order norms governing norm adoption.

The same moves, and the same limitation, arise for the parallel expressivist strategy for meeting the other condition of internal adequacy that I mentioned. This was that an adequate account of normative judgments must allow for the idea that one's normative judgments are independent of the fact that one holds them, and that therefore one's judgments might be mistaken.

According to the expressivist account I have been describing, the thought that I might be mistaken in thinking that p is a reason for me to do a could be understood as expressing my plan to count p in favor of a , in circumstances like mine, but at the same time expressing my acceptance of higher order norms governing the acceptance of such plans and the thought that these norms might turn out to support revision of my current attitude. But it is also intelligible to think that these higher order norms might themselves be mistaken. One way of expressing this would simply be through the thought that I might come, in the future, to hold different higher order norms, ones that would mandate changing my attitudes toward p and a . But this would not capture the thought that those attitudes might be mistaken, since it does not distinguish between changes that are corrections and ones that involve falling into error.

The expressivist strategy I am considering is to make this distinction by appeal to yet higher order norms that might require change in my current norms of attitude revision. Since it is intelligible to think that I might be also mistaken in accepting *these* norms, the possibility of regress looms again. But a deeper problem flows from the fact that the norms appealed to, at any level, to mark the difference between changes in attitude that are corrections and those that are errors, must, on the expressivist account, be norms that the person in question currently holds. This means

that the possibility that one might be fundamentally in error in one's normative beliefs is not intelligible on this account.¹² The account can make sense of the thought that someone else might be in fundamental normative error, however. So, as Andy Egan argues, the implication of this expressivist view seems to be that each of us must regard him or herself as uniquely immune to this possibility. This is a very odd result.

A cognitivist view of the kind I am advocating avoids these difficulties. It cuts off the regress of norms at the start by holding that when one makes a normative judgment one claims that this judgment is correct, rather than merely expressing one's acceptance of norms that support it. A cognitivist would agree that if one of one's normative judgments is mistaken, then there is a (correct) norm of attitude revision which would call, in the light of certain information that one now does not possess, for the revision of this judgment. But the thought that one's judgment might be incorrect does not involve endorsement of any particular such higher order norm.

Even if the cognitivist can avoid this regress problem, however, I have some doubts about the intelligibility of the thought that all of one's normative judgments might be mistaken, even on a cognitivist view. Might I be mistaken in thinking that pain is in general to be avoided rather than to be sought? I do not see how I could. What kind of mistake might I be making? To ask this question is to ask what there is about my current view that some norm (correct) would find faulty. It is not to endorse an particular such norm or higher order norm of norm adoption. Still, the idea that

¹² See Egan, "Quasi-Realism and Fundamental Moral Error," *Australasian Journal of Philosophy*, 2009. Egan's charge, in the last section of his paper, that Blackburn's view leads to a crude subjectivism, seems to me mistaken. I am concerned here with his argument in earlier sections.

there might be some such fault seems inconceivable. But this inconceivability is a substantive matter.

The correctness of a normative judgment, in my view, is independent of one's own view of the matter, which therefore might be mistaken. This accords well with the natural understanding of interpersonal argument and disagreement about normative matters. Since the correctness of a normative claim is independent of your opinion and of mine, it is something we can discuss and potentially disagree about, just as we can disagree about who was Prime Minister of Great Britain in 1917.

In these respects the acceptance of a normative judgment is like a belief about some empirical or mathematical question. But many object to interpreting acceptance of a normative judgments as a belief. Simon Blackburn, for example, says that this way of understanding ethical judgments in particular is “destabilized” by “questions of epistemology and of why we should be concerned about the ethical properties of things.”¹³ I will take up the question of epistemology in my next lecture. But there are several things to be said here about Blackburn's second question.

First, it makes a difference that Blackburn is speaking about *ethical* truths and beliefs, rather than about normative truths and beliefs about them. As I have said, there is an intelligible question of why we should care about the moral rightness of our actions. But this is a *normative* question, one that asks for a reason and is answered by giving one. As I pointed out in my first lecture, things are quite different when the subject is normative truths in general. There may still be an epistemological question about how we come to know such truths, but the question “Why should one

¹³ *Ruling Passions*, p. 80.

care about what reasons one has?" is nonsensical if it is understood to be asking for a reason.

Perhaps Blackburn's question "why we should care," asked about normative truth in general, is not one that asks for a reason, but is rather a question about the rational authority of reasons, a version of Korsgaard's question of how normative truths could "get a grip on" an agent. If so, this would support the suggestion that I made in my first lecture, that not only Kantians but also proponents of many other theories, including some desire theories and, in this case, expressivism, believe that the authority of reasons must be grounded in something that an agent already accepts (according to Kantians, it must be grounded "in the agent's own will," and this would also be the view of some expressivists, on the interpretation I am considering, if accepting a plan or imperative is an expression of the agent's will.)

There is a deep divide here. As I said in my first lecture, when it is true that p is a reason for someone to do a then the "grip" that p has on that person is just the normative relation of "being a reason for." Since the question of whether something is a reason is a question one asks in deciding what one's "will" in the matter is to be, grounding the answer to such questions in the agent's will does not seem like a possibility. Although it may sound excessive to say it, the claim that all claims about reasons must be so grounded threatens to eliminate reasons altogether.

A less extreme, and more common, way of understanding Blackburn's question is to take it as a question about how a (mere) belief about reasons could explain action. This is a genuine question, and I have offered an answer to it earlier in this lecture. It does follow from that answer, however, that if the acceptance of a

normative judgment is a form of belief, it differs from other beliefs—such as empirical beliefs and mathematical beliefs—in being rationally related to intentions and actions, rather than merely to other beliefs. By a rational relation I mean a connection that it is irrational to deny. Factual beliefs can have a weaker form of “rational connection with action” insofar as they are beliefs about things that provide reasons. A person who has such a belief is open to rational criticism for not treating this fact as a reason. But the failure to do so is not always irrational. By contrast, it is irrational to judge some consideration to be a reason to do a, and then refuse to treat it as a reason.

It might be maintained that if acceptance of a normative judgment has this kind of rational connection with action, then it is not a belief, since (by definition?) beliefs do not have such connections. Absent some further argument, however, this claim seems to me merely stipulative. Little turns on the term ‘belief’ as long as it is recognized that judgments about reasons can be correct or incorrect independent of their being made, and thus that they behave like beliefs in interpersonal argument and disagreement.

Arguments have, however, been offered against the possibility of a state (whether it is called a belief or not) that has the features I am claiming for the acceptance of a normative judgment. “Direction of fit” arguments maintain that no state can both have standards of correctness (a ‘mind to world’ direction of fit) and rational connections to intention and action (a “world to mind direction of fit.”)¹⁴

This claim has some plausibility as long as “the world” in question is taken to be the

¹⁴ See, for example, Michael Smith, *The Moral Problem*, pp. 112-125.

natural world of physical objects, causes and effects. For any proposition p about that world, a belief that p has a mind to world direction of fit—that is to say, a person is open to rational criticism if he or she does not modify this belief in the face of credible evidence that p is false. Any such p might also be a good reason for some action a . If so, then a person who believes p and fails to treat it as a reason for a is making a normative error and thus open to a kind of rational criticism. But this criticism is appropriate in virtue of the truth of a further normative claim $R(p, c, a)$, not simply in virtue of the fact that the agent believes p . So a belief that p is linked to standards of correctness (must “fit the world”) simply by being the kind of state that it is, but it is not rationally tied to action in this same way.

This argument depends, however, on the assumption that the belief in question is a belief about the natural world. If it is not—if the relevant standard of correctness is not “fitting with” the natural world but some other form of correctness—then the second half of the argument fails. In particular, if the belief in question is a belief that p is a good reason to do a , then it is true simply in virtue of being the kind of state that it is (and not in virtue of any further normative fact) that a person who has that belief would be irrational in refusing to treat p as such a reason. The plausibility of the argument that a state cannot have both “mind to world” and “world to mind” directions of fit is limited to cases in which “the world” referred to in both cases is the natural world. The tendency to think that this argument rules out interpreting normative judgments as a kind of belief is thus another instance of the tendency mentioned in my second lecture, to identify the set of all things independent of us about which our opinions can be correct or incorrect with “the natural world.”

Nonetheless, there are reasonable questions about the how the idea of correctness that is supposed to apply to normative judgments is to be understood. On expressivist views the essential content of such judgments is given by some “active” element, such as adopting a plan or accepting an imperative, which renders these judgments incapable of being true or false. My strategy has been to “export” this active element—to account for the distinctive practical significance of judgments about reasons by appeal to the idea of rational agency. The remaining content—the claim that something is a reason—is left as something that can be true or false, that one can be mistaken about, and that can function in interpersonal discourse like any other proposition. The question is how this residual content—the claim of correctness—is to be understood. The obvious significance of judgments about reasons lies in their rational links with action. If this is all there is, then it would seem that the “cognitivism” I am proposing will just be another form of expressivism.¹⁵

This question might be answered by a metaphysical account of the truth conditions of normative judgments. But my minimalist interpretation of normative truth rules out such an account. The idea that normative judgments are correct when they correspond to the normative facts is no explanation if these “facts” are, as I have suggested, merely “the reflection of true thoughts.” So, it may be said, the question remains what the content of these thoughts is and what makes these thoughts true?

At this point, I believe, defenders of irreducibly normative truths must dig in their heels. The idea of some consideration’s being a good reason for some action or

¹⁵ This would be the analog for my view of the charge that Blackburn’s quasi-realism cannot have the virtues he claims for it without becoming a form of realism. See Gideon Rosen, “Blackburn’s *Essays in Quasi-Realism*,” *Nous* 32 (1998) pp. 386-405.

attitude is a perfectly intelligible one. To believe that some consideration is a reason is not the same thing as treating that consideration as a reason in subsequent deliberation—there is such a thing as irrationally failing to act in accord with the reasons one believes oneself to have. Given the intelligibility of this idea, and the fact that taking it at face value provides the best fit with our practices of thinking about reasons and arguing about them with others, we should reject it only if it gives rise to some difficulties that cannot be answered satisfactorily. I have argued that the idea of irreducibly normative truths does not have implausible metaphysical implications, and that the connection between beliefs about reasons and subsequent action can be satisfactorily explained.

There remains the “epistemological” question of how we can know what reasons we have. The question “what makes normative judgments true?” might be understood as a way of asking this question, that is to say, a way of asking how the correctness of normative judgments can be established. I will take up this question in my next lecture, and as I will say then, I believe that there are grounds for limited expectations on this score. If one were to have a systematic account of the procedures through which normative truths can be established, then one might simply identify the idea of a normative judgment’s being correct with its being established in this way. For various reasons, I doubt that this strategy will work, but I will explore the possibility in my next lecture.

Aside from worries about how the idea of the correctness of normative judgments is to be understood, however, questions may also be raised about the importance that idea should have for us. I have suggested that it is important in two

contexts: in making sense of the idea that the correctness or incorrectness of our normative judgments is independent of our making those judgments, and in interpreting interpersonal discourse and disagreement about normative questions. Each of these forms of importance may be questioned.

Suppose that you and I disagree about whether the fact that someone injured me is good reason for me to injure him in return. Perhaps I maintain that it is, and you deny this. Suppose we go on for some time, arguing about this and adducing all the considerations that either of us can think of to get the other to change his mind, but that we still disagree. It would be pointless and empty for me, at this point, to insist, as if it were a trump card, “But my view is *correct*. It *is* a reason!” Such an appeal to “correctness” would be mere foot-stomping.¹⁶

Similarly, in the intrapersonal case, if one believes that something is a reason, it is natural to think that it would be a reason whether or not one believed that it was. But why should this be so important to us? If we are convinced that something is a reason, and ready to act on it, why should we be concerned to have the imprimatur of some independent standard of correctness? As Nietzsche would say, the need for the prop of such a standard betrays a kind of weakness.¹⁷ In a similar vein, Simon Blackburn says it is sad, that some people should feel this need for Apollonian authority, rather than being content to accept the motivation provided their own, contingent, emotions and desires.¹⁸

¹⁶ As Christine Korsgaard observes. See *The Sources of Normativity*, p.

¹⁷ See *Genealogy of Morals*, Third Essay, esp. sections 26-28.

¹⁸ *Ruling Passions*, pp. 88-91. See also the concluding paragraphs of “Am I right?” *The New York Times Book Review*, February 21, 1999, p. 24.

As I will say in my next lecture, I think there are cases in which the interpersonal version of this worry points toward a genuine issue. But neither worry provides grounds for rejecting concern with the idea that normative judgments can be correct or incorrect. To insist, at the conclusion of an unsuccessful attempt to persuade someone, that your normative judgment is correct, is indeed unhelpful foot stomping. But this is equally true when the disagreement is about some matter of empirical fact. So this does not show that the idea of correctness is misplaced, or better done away with.

Moreover, the idea that when we disagree about a normative question there is some fact of the matter we are disagreeing about, independent of each of us and which neither has any special authority to determine, provides if anything a more attractive picture of the situation than the idea that we are each simply trying to get the other to adopt the same plan that we have adopted. Of course, the fact that this interpretation of disagreement is attractive, if it is, is not any reason to think that it is true—that normative judgments can be correct or incorrect. My aim in mentioning the attractiveness of the idea is just to rebut the suggestion that to be concerned that there be such a notion of correctness betrays a desire to claim an implausible and unattractive kind of authority for one's position.

Nor, when several alternative courses of action seem appealing, does it indicate a kind of weakness to ask oneself which one really has more reason to do. The idea that it does may derive from thinking that being concerned with the correctness of one's normative beliefs involves looking for some authoritative standard "outside of ourselves" that will tell us what we ought to do. But this idea of

an “outside standard” is a misleading metaphor. For any “outside standard,” in the form of a set principles or precepts, there is the question, “Why do that?” But when I arrive at a conclusion about the correctness of a normative judgment—that I really do have reason to do what will save my life, or to avoid pain—there is no such further question. These conclusions carry their own normative authority, as it were. They do not need to derive it from some further source. The question of correctness is the question of whether they do have this authority—whether the considerations in question are really reasons. There is no further question beyond this one. But there are serious questions about how it can be answered, and I will turn to these in my next lecture.