

Controlling Attitudes

Pamela Hieronymi
hieronymi@ucla.edu
October 12, 2004

I hope to show that, although belief is subject to two quite robust forms of agency, “believing at will” is impossible; one cannot believe in the way one ordinarily acts. Further, the same is true of intention: although intention is subject to two quite robust forms of agency, the features of belief that render believing less than voluntary are present for intention, as well. It turns out, perhaps surprisingly, that you can no more intend at will than believe at will.

I hope to show that believing could not be “voluntary,” that is, one could not believe in the way one can perform an ordinary intentional action. You could not believe in the way that you can raise your right hand or look left or prepare dinner. Moreover, this is, as suspected by Bernard Williams, a conceptual matter, traceable in part (but only in part) to the relationship between belief and truth: beliefs “aim at truth” or “purport to represent reality.” However, I also hope to show that, while there is a sense in which belief is not “voluntary,” it is nonetheless subject to two quite robust forms of agency, and, further, that these two forms of agency are also exercised with respect to intention. In fact, we will see that the features of belief that render believing less than voluntary are present for intention, as well—even without the aim at truth. It will turn out, perhaps surprisingly, that you can no more intend at will than believe at will.

I. BELIEVING VOLUNTARILY

In one sense, it can seem obvious that we can’t “decide to believe,” or believe at will, in the way we can decide to act, or act at will. While any able-bodied person can, at will, raise her right hand, look to the left, or turn herself about, it seems that a well-functioning believer cannot, in the same way, just decide to believe that it’s raining outside, or that it’s Christmas Day, or that Clarke Gable is the President of the United States. You don’t have the kind of discretion over what you believe that you generally have over what you

do. You can't just believe whatever suits your fancy. However, in another sense of the phrase, no one would doubt our ability to believe "by deciding:" we form at least some of our beliefs by making up our minds about what is true. We often decide, e.g., whether we can make it to the movie on time, whether a product is fairly priced, or whether it is likely to rain. Believing sometimes involves deliberation, or the drawing of a conclusion; in such cases, believing does not seem to be an involuntary experience, like sneezing.

In his "Deciding to Believe," Bernard Williams suggests a possible explanation for why it is that you can't believe in the way you can look left or raise your right hand: beliefs necessarily "aim at truth." He provides the following oft-quoted argument:

One reason [I cannot believe at will] is connected with the characteristic of beliefs that they aim at truth. If I could acquire a belief at will, I could acquire it whether it was true or not. If in full consciousness I could will to acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e., as something purporting to represent reality. At the very least, there must be a restriction on what is the case after the event; since I could not then, in full consciousness, regard this as a belief of mine, i.e., as something I take to be true, and also know that I acquired it at will. With regard to no belief could I know . . . that I had acquired it at will. But if I can acquire beliefs at will, I must know I am able to do this; and could I know that I was capable of this feat, if with regard to every feat of this kind I had performed I necessarily had to believe that it had not taken place?¹

Because beliefs aim at truth, you could not regard anything that you created in yourself without regard for truth as a belief, and so you could not believe at will. Because beliefs purport to represent reality, believing is beyond our voluntary control.²

Notice that aiming at truth, or purporting to represent reality, is a necessary feature of any belief. An attitude that didn't "aim at truth" would not qualify as a belief. If this necessary feature is what accounts for our inability to believe at will, then believing would be less than voluntary, it seems, for any believer. Our inability to believe at will thus contrasts, Williams thinks, with our inability to blush at will. While we cannot, in fact, blush at will, we can readily imagine ourselves able to do so—able to decide to

blush just as easily as we can decide to raise our right hand. It does not seem that blushing must be involuntary for any blusher.

In his “Why Is Belief Involuntary?” Jonathan Bennett considers Williams’ argument and provides a counterexample to the claim that, because beliefs aim at truth, believing is not voluntary for any believer. He does so by evoking a fictional community:

Credam is a community whose members can [decide to believe]. It doesn’t happen often, because they don’t often think: “I don’t believe that p, but it would be good if I did.” Still, such thoughts come to them occasionally, and on some of those occasions the person succumbs to temptation and wills himself to have the desired belief. . . . When a Credamite gets a belief in this way, he forgets that this is how he came by it. The belief is always one that he has entertained and has thought to have some evidence in its favor; though in the past he has rated the counter-evidence more highly, he could sanely have inclined the other way. . . . The trick cannot be worked if the protective forgetfulness would require that the rest of the person’s beliefs be drastically rearranged. . . . After successfully willing himself to have a certain belief, a Credamite may later get evidence that that is what he has done; e.g., someone may tell him. Then he either rejects the evidence or loses the belief. . . . So each Credamite knows that he sometimes wills himself to believe something, even though it is never true that he *now* has a belief which he *now* remembers having willed himself to acquire.³

Bennett’s example seems to me a successful counter to Williams’ argument, as given.

The Credamites are capable of believing “just like that,” and yet they can regard the beliefs at which they thus arrive as aiming at truth.⁴

Despite his own counter-example, Bennett thinks Williams’ intended point is right—our inability to believe by just deciding to, unlike our inability to blush by just deciding to, is not merely a contingent limitation of human psychology. Bennett’s paper was originally intended to counter his own counterexample and offer an alternative argument to Williams’ conclusion. Unfortunately, Bennett opens the paper by admitting its failure.

I agree with Bennett that Williams’ intended point is right. Despite the plausibility of the Credamites, believing in the way you raise your right hand or look left is a conceptual impossibility. Further, this is due in part to the fact that beliefs purport to represent

reality, or “aim at truth.” I hope here to provide a successful alternative explanation of why we can’t believe in the way we can act.

Approaching this problem will require examining what it is for an activity to be “voluntary,” or done at will, or performed as an ordinary action.

Bennett suggests an account of voluntariness that will provide a starting point. According to Bennett, voluntariness is “responsiveness to practical reasons. . . . Actions are voluntary in that sense, and beliefs seem not to be.”⁵ By “practical reasons,” Bennett means reasons that bear on the question of what *to make* true, as opposed to reasons that bear on what *is* true.⁶ If you take yourself to have reason to (make it true that you) raise your right hand, then you can, Bennett thinks, immediately, for those reasons, raise your right hand. To believe voluntarily, for Bennett, would be to take yourself to have reason to make it true that you believe *p*, and then be able, immediately, for those reasons, to believe *p*. The Credamites seem to believe voluntarily, in this sense: when they have reason to make it true that they believe, it seems that they can, immediately, for these reasons, believe. So according to Bennett, an activity is voluntary if it can be done immediately in response to practical reasons.⁷

The qualification “immediately” is important. As Bennett acknowledges, sometimes even we ordinary mortals can, in response to practical reasons, undertake some process by which we might *bring ourselves* to believe this or that. If I have reasons for making it the case that I believe the lights are on in my office, I can, for those reasons, get up and throw the switch, thus bringing myself to believe.⁸ More dramatically, philosophers have imagined science-fictional pills that can induce both belief and amnesia about the means of belief acquisition. By taking such a pill, I might bring myself to believe this or that.

Bennett calls any such process believing “mediately:” for practical reasons you are motivated to do something *else*, which brings you to belief. Bringing yourself to belief by any such mediating process is not believing at will. Rather, to believe at will, it seems one must believe *immediately* for practical reasons—one must decide to believe, for such reasons, and thereby believe, without having to do anything else to produce that belief. We can, in this sense, act at will: we can decide to act, for practical reasons, and then immediately act, without having to do anything else to bring ourselves to act. In fact, it seems we *must* be able to act immediately, without having to do anything else to bring ourselves to act, on pain of infinite regress. Unlike us, the Credamites seem thus able to believe.

Notice, though, that the sense of “immediacy” at issue here is not at all clear. First, not just any believing “immediately in response to practical reasons” will count as believing *at will*. By threatening me with grave harm unless I believe *p*, you might cause me to believe *p*, immediately—perhaps, faced with such frightening threats, I just find myself believing. Though in some sense an “immediate” response to practical reasons, this is not believing at will.⁹

Bennett seems to identify the immediacy at issue with the immediacy of “basic action.” Basic actions are “immediate” in that they can be performed without having to do anything else to perform them. Raising one’s arm is thought to be a basic action, for able-bodied humans; preparing dinner is not.¹⁰

But note how odd it would be to use *this* sense of immediacy to distinguish what ordinary humans do, in bringing themselves to believe through some process, from believing at will—believing as an ordinary action. It would suggest that bringing

yourself to believe by the execution of some clunky, multi-step process does not qualify as believing voluntarily simply because of the clunky process. But surely whatever sense of “voluntary” divides believing from raising one’s right hand also divides believing from preparing dinner. The fact that I cannot prepare dinner without chopping the vegetables, turning on the burner, and heating the oil does not render my action any less “voluntary,” in the relevant sense. It simply renders it non-basic. Likewise, if I were somehow so divinely constituted as to be able to make it the case that dinner is prepared simply by deciding that it be so, my action would be no more voluntary, in the relevant sense. It would simply be much easier.

But if the complexity of an action like preparing dinner does now show that one cannot prepare dinner “at will,” in the relevant sense, it is unclear why the need to engage in some complex process in order to believe in response to practical reasons should show that one cannot believe “at will.” Or, if we agree that what one does in bringing oneself to believe for practical reasons is *not* believing at will, then it is unclear why shortening the process—making it ever easier, until we reach the limiting case in which we are equipped with a psychological mechanism by which we can bring ourselves to believe by deciding to—should somehow transform something that was not a voluntary action into a basic one. Yet it seems the Credamites believe at will; if they do not, we need an account of why they do not.

We now have our task set out for us: show why one cannot believe in response to practical reasons in whatever sense one can ordinarily act in response to such reasons. Show why believing cannot be an ordinary intentional action.

To address this question one must, of course, examine a large range of philosophically difficult topics: beliefs, actions, intentions, the reasons for them, and the agency we exercise with respect to them. My strategy will be as follows. I will start by considering belief, reasons for believing, and agency with respect to belief. I will then consider the same set of issues for intention. In the fourth section I will sketch an account of intentional action before returning, in the fifth section, to believing at will.

II. BELIEVING

Let's start by examining belief—in particular, the relation between one's beliefs and one's assessment of what's true. Most basically, it seems a truism that one believes *p* insofar as one takes *p* to be true. However, this formulation has been questioned. Upon examination, it fails to distinguish believing *p* from supposing *p* (for the sake of argument, perhaps) or imagining *p*.¹¹ After all, to suppose *p*, for the sake of argument, is equally to take *p* to be true, for the sake of argument. I hope to avoid this difficulty by pointing out that, if one believes *p*, one is thereby vulnerable to criticism or critique under certain standards—the standards of justification, warrant, or consistency that govern belief. In self-reflective, language-using creatures like us, believing *p* will thus leave the believer answerable to requests for a particular kind of justification—a justification that bears on *p*'s truth—and open to possible charges of inconsistency. If you believe *p*, you can be asked why you do—what you take to show it to be true—and you can be asked to explain how that belief comports with your other attitudes and even your actions. The kind of answers you must give, if you are to be justified and consistent in believing, are quite different than the answers you must give to be justified and consistent in supposing or imagining.¹² If one didn't understand the (in principle) applicability of the appropriate

set of questions or criticisms, that would show that one doesn't understand what it is to believe.¹³ Thus it should be a truism that you believe *p* just in case you take *p* to be true in a way that leaves you answerable to certain questions and criticisms—namely, those that come with believing *p*. I will mark this by saying that you believe *p* just in case you are “committed to *p* as true,” where to be committed to *p* as true is to take *p* to be true in a way that leaves you thus vulnerable. Belief is, I will say, “commitment-constituted.”¹⁴

In sum, I take the thought that belief “aims at truth” or “purports to represent reality” to mark that fact that to believe that *p* is to be committed to *p* as true—to take *p* to be true in a way that leaves one answerable to certain questions and criticisms.

Reasons for Believing

With this understanding of belief, we can turn to reasons for believing and notice an important ambiguity in the idea of a reason for belief. I take a reason to be a consideration, i.e., some fact or proposition, that bears on a question.¹⁵ The fact that the butler wanted revenge bears on the question of whether he is guilty, and so is a reason (though not a conclusive reason) for believing he is. Notice, though, that there are two distinct questions on which a consideration can bear, and so count in favor of believing *p*. A consideration can count in favor of believing *p* either by bearing on the question of whether *p*, as we have seen, or, instead, by bearing on the question of whether the belief that *p* would be good, useful, appropriate, desirable, important, etc., to have. If the only way to save my life is to believe that the butler is guilty, then the fact that believing he is guilty would save my life bears on the question of whether it would be good for me to believe, and so counts in favor of believing.

Reasons can count in favor of believing in these two quite different ways because one's belief is not only one's take *on* the world—one's assessment of what is true—it is also a fact *in* the world—a piece of one's psychology, a state of mind. As self-reflective creatures, we are capable of thinking about the fact that we believe this or that, and what we believe is often of some importance to us. Thus a consideration can count in favor of believing not only by showing that the content of the belief is true but also by showing the belief, as a piece of psychology, is good to have.¹⁶ Christian Pillar has labeled this distinction as between “content-related” and “attitude-related” reasons for belief.¹⁷

The content-/attitude-related distinction should, I claim, be drawn as I have drawn it—by considering the question on which the consideration is taken to bear. Notice that the question on which an attitude-related reason bears—the question of whether the belief is good to have—must mention the belief of which the reason counts in favor. In contrast, content-related reasons manage to count in favor of a belief by bearing on a question that typically doesn't mention the belief. If you believe that the butler did it, your content-related reasons for this belief might include his motives for revenge, his ready access to the home, and his long-standing fascination with toxic substances. These considerations bear on the question of whether, in fact, the butler did it. That question does not make any mention of your beliefs. Thus drawn, the content-/attitude-related distinction marks whether appreciating the reason requires a certain sort of reflection.

Drawn in this way, the distinction between content- and attitude-related reasons is not, in principle, an exclusive one; the answer to one of these questions can bear on the other question, and so reasons that provide an answer to one question can, in turn, become reasons of the other sort. For example, quite often (though not always), it is

good to have true beliefs. Thus, by showing the content of a belief to be true, content-related reasons might also thereby show the belief is good to have.¹⁸ For self-referential beliefs (such as the admittedly odd, “This belief is good to have”), attitude-related reasons, which show the belief good to have, may also, thereby, bear on whether its content is true.¹⁹

To discuss agency over belief, we want a slightly different distinction, one which is drawn by considering the relationship between appreciating a reason and believing. Suppose you take certain considerations to bear on whether p , and, further, you take them to be sufficient to show that p . By this, I do not mean merely that you take yourself to have reason enough that, if you were to believe p , you could not be criticized as unreasonable or unjustified for doing so. I mean something stronger: you yourself find the reasons convincing, you are convinced by them.²⁰ Of course, if you take certain reasons to show that p , you therein believe p .²¹ Thus the reasons taken to bear positively on whether p —those taken to be content-related reasons for the belief that p —are also what I will call “constitutive reasons” for the belief that p . They support the commitment constitutive of the belief. By finding such reasons convincing, you therein believe.²²

Consider, now, the *remaining* reasons for believing p —those taken to count in favor of believing p independently of whether p . Finding these reasons convincing does not amount to believing p . Suppose that your life depends on your ability to give convincing testimony that the butler is innocent, but you do not believe he is innocent, and you are a terrible liar. Now, you might think that you have overwhelming reason showing it good to believe he is innocent—it would save your life. Yet, the fact that you take this reason to be convincing does not show that you believe him innocent. Instead, it shows that you

have a *second-order* belief *about* the belief that the butler is innocent: you believe it would be good to believe him innocent. Since finding these reasons convincing implies nothing about whether you have the belief of which they count in favor, I call them “extrinsic” reasons for that belief.

The constitutive/extrinsic reason distinction, drawn by considering the relation between finding certain reasons convincing and believing, is closely related to the distinction between content- and attitude-related reasons. The class of constitutive reasons will be identical to the class of content-related reasons—including those attitude-related reasons which, by showing the belief good to have, also thereby bear on the truth of its content. Extrinsic reasons are simply those that are not constitutive. As it turns out, the extrinsic reasons will be identical to the remaining attitude-related reasons: they count in favor of the belief by bearing on whether the belief is in some way good to have. I see no other question on which a reason could bear and so count in favor of the belief.²³

Finally, notice that a fact or consideration becomes a reason only as it bears on some question. A particular consideration might be taken to bear independently on the content- and attitude-related questions, and so provide both a constitutive and an extrinsic reason for the same belief. Suppose, e.g., I have an exam next Wednesday. Now, I might take the fact that today is Friday to count in favor of believing that I will do well on the exam, both by bearing on the likelihood that I will do well (because the exam isn’t until Wednesday, and so I have plenty of time to study) and by bearing, independently, on the question of whether it would be useful to have the belief that I will do well (because I want to have fun over the weekend, and I won’t if I am worrying over the exam). Here, I take the same consideration, that today is Friday, to bear on two separate questions. It

shows both that it's likely that I will do well and, independently, that it would be beneficial to believe that I will do well. Thus the same consideration provides me with two reasons for believing that I will do well, one constitutive, the other extrinsic.²⁴

Bennett defined “practical reasons” as reasons which bear on what to make true, as opposed to reasons that bear on what is true. Presumably, reasons which bear on whether to make it the case that you believe do so by showing something good, in some way, about believing, without showing the belief true. So presumably practical reasons for a belief are extrinsic reasons for that belief. To believe at will, then, one would have to be able to believe for *extrinsic* reasons.

Controlling Belief

We have so far established that to believe *p* is to take *p* to be true in a way that leaves one answerable to certain questions and criticisms, we have drawn a distinction between constitutive and extrinsic reasons for believing, and have determined that to be able to believe at will one must be able to believe for extrinsic reasons. I next want to introduce a distinction between two kinds of agency, or control, that we can exercise with respect to our own beliefs.²⁵

Consider, first, our control over ordinary objects like cars and coffee cups. We exercise this sort of control when we manipulate some ordinary object to accord with our thoughts about it. We typically control ordinary objects by performing intentional actions which affect that object in the way we intend. Of course, our control over such objects is never unlimited. We are subject to physical and temporal limitations, to limitations of skill, and to luck. Importantly, the degree to which we exercise control

over some object is measured not by the absence of such limitations—as though we would exercise greater control over our coffee cups if they did not obey the laws of physics—but rather by our ability to navigate, manage, and make use of those limitations in order to accomplish our purposes. In fact, in many cases, exactly those features that seem to limit our control also make such control possible. We can control ordinary objects at all only because we know they observe certain regularities. I will call this first sort of control “manipulative” or “managerial” control.

Consider, next, our beliefs. Notice that our own beliefs also accord with our thoughts, though not in the same way that ordinary objects might. Consider a belief that it will take forty-five minutes to drive to the airport. If we learn we must leave at rush hour, we then reconsider how long the drive will take, and our belief changes accordingly. So, in a certain sense, we might also be said to “control” our own beliefs. While we control our cars and cups by manipulating them according to our purposes, we control our beliefs by evaluating (and reevaluating) what is true. Call this “evaluative” control.²⁶

Of course, the evaluative control we exercise over our beliefs differs sharply from the ordinary control we exercise over ordinary objects. Our control over our beliefs is *immediate* in a way that our control over ordinary objects is not. When you change your mind about what is true, you have therein, *ipso facto*, changed your beliefs. In contrast, when you change your mind about which shoes to wear, you haven’t therein changed your shoes. You have only changed your mind. To restate: By answering positively the question of whether *p*, you have, therein, believed *p*. By answering positively the question of whether to change your shoes, you have, therein, *intended to change* your

shoes—but you haven’t yet changed them. Even in cases of basic action, such as raising your right hand, there remains a conceptual difference between deciding to act (intending) and acting. There are possible cases in which one decides to act but is prevented from acting (because, say, your arm suddenly goes numb). In contrast, there are no possible cases in which you answer positively the question whether *p* but are prevented from believing *p*. In answering the question positively, one has already, therein, believed. The immediacy of evaluative control is thus not temporal or causal but rather a consequence of the constitutive relation between the commitment to *p* as true and the belief. Since we believe *p* insofar as we are committed to *p* as true, when we take *p* to be true in the relevant way, our beliefs therein change.²⁷

Further, while immediate, our control over our beliefs is typically *indirect* in a way our control over ordinary objects often is not. When controlling your cup or your shoe, you have typically decided to do something to or with *it*. In contrast, you control your belief by answering a question about its content—a question that only rarely concerns your own psychology.²⁸

Managing and Manipulating Beliefs

Can we exercise manipulative control over our beliefs? Can we manage our beliefs in something like the way we manipulate our chalk and coffee cups? Given that we are reflective creatures, able to think about our own beliefs, and given that our beliefs interact more-or-less reliably with other features of the world, it would be quite surprising if we couldn’t turn our manipulative, managerial energies toward them. (After all, given the right circumstances we can manage the beliefs of others quite effectively.) In this section I hope to show not only that we can manage our beliefs in something like the way we

manipulate our cars and coffee cups, but that doing so is familiar to us. While the commitment-constitution of belief does impose a constraint on our attempts at belief management—we cannot produce a belief in ourselves without coming to be committed to its content as true—I will argue that this constraint does not impugn the claim that we exercise managerial control over our beliefs. In fact, some such constraint is necessary for control.

Suppose you can't fall asleep because you are worried about whether your friends arrived safely home through the storm. Wanting sleep, you have an extrinsic reason for the belief that they are safely home. This extrinsic reason gives you reason to produce in yourself a belief. The obvious thing to do, in this case, is to conduct a little investigation: call your friends. If you find them home, you will have brought yourself to the desired belief.²⁹

Notice that the reasons you take yourself to have in the process of conducting this investigation will observe a particular division of labor, due to their distinct relations of justification to the desired belief. Reasons which merely show the belief good to have—extrinsic reasons—support the undertaking and continuance of the investigation. But since they do not bear on the question of whether the belief is true, they cannot legitimately support any particular conclusion of investigation. The constitutive reasons unearthed by the investigation should do that work. Thus the two sorts of reason, because of their differing relations of justification to the belief, support two coordinated but distinguishable activities: conducting an investigation and arriving at a particular conclusion. (The same division of labor appears if we bring ourselves to a desired belief

not by conducting an investigation but instead by changing the world to make the belief obviously true.)

So, we can often manage ourselves into a belief by bringing ourselves, by some process that provides us with what we take to be good constitutive reasons, to be committed to the content of the belief as true.³⁰ But suppose no such reasons are forthcoming. Can you, in such a hard case, manipulate yourself into a belief?³¹

As already mentioned, philosophers have imagined various science-fictional ways to produce belief: pills, surgery, special psychological mechanisms. By such means, you somehow simply make it the case that you are committed to *p* as true—you somehow simply make it the case that you take *p* to be true in such a way that you are now answerable to all the questions and criticisms that come with believing *p*. It seems imaginable that, given the right sort of circumstances, you could bring yourself to incur such a commitment by some such process. I don't claim that it is impossible.³²

If possible, such a process might simply produce irrationality: it might simply saddle you with conflicting commitments—i.e., with a belief that you also find (perhaps in your better moments) implausible. Such irrationality certainly seems possible. While it is necessarily both incoherent and unstable, it nonetheless seems sometimes to actually obtain. Thus I can provide no reason, in principle, to think that we couldn't devise a way to bring ourselves into it. The limitations on self-induced irrationality seem merely “technological,” not conceptual.

Alternatively, a science-fictional process for arriving at a belief might avoid producing irrationality by somehow ensuring that you do not immediately encounter other evidence or remember other beliefs that you would take to show the new belief

false—including beliefs about how you acquired the new belief. As long as we encounter no reason to doubt a belief, we can continue in it, without irrationality.³³

So, the commitment constitution of belief does constrain our ability to manage our own beliefs. Yet we have just seen that this constraint does not eliminate the possibility of belief management (even in those cases in which good constitutive reasons are lacking), and so does not falsify the claim that we can exercise managerial control over our beliefs. In fact, we can design to produce a belief in ourselves precisely because we know that a belief is constituted by commitment to its truth.

We have now seen two distinct forms of agency we are able to exercise with respect to our own beliefs. Notice, further, that these two forms of agency seem to align with the two kinds of reasons we might take to count in favor of believing: we seem to exercise evaluative control in response to constitutive reasons and managerial control in response to extrinsic reasons.

III. INTENDING

Before returning to our challenge it will be useful both to consider, briefly, the parallel claims for intention and to provide a sketch of intentional action.

Intentions, like beliefs, involve a certain sort of answerability. It is now quite standard, in philosophy of action, to think of intending as settling the question of what one will do. Having settled that question, like having settled a question regarding what is true, leaves one open to certain questions and criticisms. If I intend to do some grading tonight, or to clean my apartment, then I am committed to grading or to cleaning in a way that I am not if I merely wish to, hope to, or predict that I will. I can be asked why I think

□-ing is the thing for me to do, and my intention to □ can be inconsistent, in familiar ways, with both my other attitudes and my actions.³⁴ Further, it seems that I am rightly subject to these particular requests for justification and charges of inconsistency only insofar as I intend to □. Thus an intention, like a belief, seems at least in part commitment-constituted. An intention is a commitment to doing something, where a belief is a commitment to a claim as true.³⁵

I mean to use “intention” more broadly than it is sometimes used. I do not mean for intentions to be restricted to future actions or to involve explicit acts of formation, but to include what are sometimes called “intentions in acting.” On this broad understanding, we can posit an intention any time a person acts intentionally. This is because, any time one acts intentionally, one is open to these questions of justification and charges of inconsistency.³⁶ By acting intentionally, it seems one has in some sense answered (however implicitly, unreflectively, unconsciously, or spontaneously) the question of whether so to act, and so has intended. (One may or may not have answered this question for reasons.)

Next, notice that we can draw the same two distinctions in reasons for intending that we have drawn in reasons for believing. We can distinguish, first, between content-related reasons (which bear on the question of whether to □) and attitude-related reasons (which bear on the question of whether it would be good, useful, desirable, appropriate, or important in some way to intend to □). As with belief, the content-related reasons for an intention support the commitment that constitutes it. We can call these the constitutive reasons for intending. By taking them to settle the question of whether to □, you will,

therein, intend to $\square$. One evaluatively controls an intention to $\square$ by answering for oneself (however implicitly) the question of whether to $\square$.³⁷

One may wonder whether there are any extrinsic reasons for an intention—reasons that show the intention good to have independently bearing on whether to act. It may seem there are not. After all, while the question of whether p only very rarely takes into account whether the belief that p is good to have, the question of whether to $\square$ readily takes into account whether or not the intention to $\square$ is good to have. One can decide to $\square$ in order to intend to $\square$. Suppose, for example, my boss wants me to intend to have the job done by the end of the day; moreover, she doesn't care about the job, but only cares about my willingness to form the intention. Now, assuming I have reason to please my boss, I have a reason for having this intention, and, importantly, the fact that the intention would be good to have can itself serve as a reason to do the job. I can decide to $\square$ simply in order to intend to $\square$. In this way, reasons which show an intention to $\square$ good to have can sometimes bear on the question of whether to $\square$. Sometimes they will be reason enough to $\square$. That is to say, attitude-related reasons for an intention are readily also *constitutive* reasons for intending. So it may seem there are no extrinsic reasons.³⁸

Nonetheless, in certain cases one can have extrinsic reasons for intending—reasons which show the intention good to have independently of their bearing on whether to act. You might, e.g., find yourself convinced that an intention to retaliate is good to have—for its deterrent effects—without thinking you thereby have reasons that show that retaliating is to be done. In such a case, you might exercise managerial control over your intention. You might take steps designed to bring yourself to intend. (You might, e.g., create for

yourself reasons to retaliate by rigging your own bombs so that, if you are attacked, they will explode on your own soil if not launched at the attacker.)

So intentions, like beliefs, seem subject to both evaluative and managerial control. One exercises evaluative control over an intention to φ by answering the question of whether to φ . If one has reasons that show the intention good to have that are not reasons to act, one might then engage in managerial control and bring oneself to intend.

IV. ACTING

Given this understanding of intention, we can sketch an account of how it is that we act in response to practical reasons. My story, which I hope relatively uncontroversial, is this: In response to reasons which bear on the question of whether to φ , you might answer positively the question of whether to φ (however implicitly, unreflectively, unconsciously, or spontaneously) and so therein (by an exercise of evaluative control), intend to φ . If nothing interferes, if you do not change your mind, and if you are not akratic, you will (at the appropriate time) execute that intention and φ (where executing the intention just is φ -ing, not some further action or activity). That is to say, you act intentionally in response to practical reasons by taking those reasons to answer the question of whether so to act, therein intending to act, and so, if all goes well, acting.³⁹

Notice that intentional action, like belief and intention, involves a kind of answerability.

G. E. M. Anscombe astutely observed that whenever one acts intentionally one is answerable to a certain sort of “why” question.⁴⁰ If you are intentionally pumping water to the house, you can be asked “why are you pumping water to the house?” where this question looks for your reasons for pumping. Importantly, Anscombe also noted that the

why-question to which one is answerable, in performing an intentional action, is sensitive to the description under which the action is intentionally performed—it is sensitive, one might say, to the intention with which one acts. The relevant why-question, if one acted with the intention to φ , is, “why did you φ ?” Why-questions that get the intention wrong (why are you poisoning the guests?) are to be rejected rather than answered (I didn’t know that I was).

The account of intending and acting I have given would lead one to expect that Anscombe’s question is, as she puts it, “given application” for any intentional action. On the given account, by acting intentionally one shows oneself committed to an answer to the question of whether to φ . The reasons which bear on this question, whether to φ are just the reasons that would answer the Anscombean question, why did you φ ? Since both the intention and the action intended involve answerability to the same set of reasons, whenever one acts intentionally one is, in a sense, in a position to answer the question, “why did you perform that action?”⁴¹

V. BELIEVING AT WILL

Let’s now return, at last, to our challenge: why can’t we believe at will? Why can’t we believe in response to practical reasons in the way we act in response to practical reasons?

Bennett identified voluntariness as responsiveness to practical reasons and identified the ability to believe at will with the ability to believe *immediately* in response to practical reasons. We have seen that the practical reasons for which one must believe, if one is to believe at will, are *extrinsic* reasons for believing.

I will now argue that believing for extrinsic reasons will always require acting upon yourself, in a certain sense, to make yourself believe. Since we do not ordinarily have to act upon ourselves, in this sense, to make ourselves act—to raise our right hand or to prepare dinner—this will show that we cannot believe “immediately” in response to practical reasons in the way we act “immediately” in response to such reasons. The immediacy at issue is the lack of managerial control. My argument will come in two stages. First, I will show that we cannot believe for extrinsic reasons by an exercise of *evaluative* control. I will then argue that any other way of believing actively in response to such reasons will require acting upon oneself.

If one could believe in response to extrinsic reasons by an exercise of evaluative control—simply by answering a question for oneself—then it would seem that one had believed at will.⁴² However, I will now argue that it is conceptually impossible to exercise evaluative control over a belief for extrinsic reasons.

To see clearly why this is so, recall, first, that a reason is a consideration that bears on a question. In taking a consideration to be a reason (however implicitly or unreflectively), one has already determined the question on which that reason bears. If I take the fact that today is Friday to be a reason to believe I will do well on the exam, I must already have in mind whether this fact counts in favor of the belief by showing that I will do well or rather only by showing that it is useful for me to believe I will do well.

Further, we have also seen that, when one answers a question for oneself (again, however implicitly or unreflectively), one might therein, *ipso facto*, arrive at a belief or an intention. If I positively answer for myself the question of whether *p*, I will therein believe *p*. If I positively answer for myself the question of whether to $\square$, I will therein

intend to $\square$. We have called arriving at an attitude by answering a question exercising evaluative control over that attitude.

Notice that there is a conceptual connection between the question answered and the attitude therein formed. If one answers a question about whether p , one will therein evaluatively control a belief that p . If one answers a question about whether it is good to believe p , one will evaluatively control a *second-order* belief about whether it is good to believe p . If one answers a question about whether to $\square$, one will therein evaluatively control an intention to $\square$. Which commitment-constituted attitude one evaluatively controls is determined by the question one answers for oneself.

Since the reasons will already specify the question under consideration, and since the question determines which attitude will be immediately formed or modified by evaluative control, the agent does not, in exercising evaluative control, have discretion over which attitude is controlled in response to which reasons. Rather, in taking a consideration to be a reason, she has already determined which attitude she will evaluatively control in response to it.

Since one evaluatively controls a belief that p by answering the question of whether p , and since extrinsic reasons are those reasons that are *not* taken to bear on the question of whether p , one cannot evaluatively control a belief that p in response to extrinsic reasons. One might instead immediately form or modify, by evaluative control, in response extrinsic reasons for a belief that p , some attitude about that belief—a second-order belief that the belief that p is good to have or an intention to bring it about that one believes p (or perhaps a desire to have the belief that p). But one cannot, as a conceptual

matter, arrive at a belief that p simply by finding convincing the reasons that you take to show only that the belief is good to have.

Thus, neither we nor the Credamites could believe in response to extrinsic reasons by evaluative control. Rather, the attitude one will form, by finding such reasons convincing, will be some attitude about the belief—perhaps an intention regarding that belief. This inability is not, of course, any shortcoming in our agency. It is merely a fact about these attitudes. One can readily form attitudes in response to practical reasons for believing p ; the attitudes simply won't qualify as believing that p .⁴³

Notice, next, that any other exercise of control over a belief must be mediated through an intention. If a person has formed a belief, in response to extrinsic reasons, without intending to, then (given the broad sense of intention at work here) she did not arrive at the belief intentionally. We also know she did not arrive at the belief by evaluative control. But, if she did not arrive at the belief by an exercise of evaluative control and she did not arrive at it intentionally, then it seems she did not arrive at the belief by an exercise of control of any kind. Rather, she was in some way simply caused to believe (perhaps by a strong desire for the belief, or a powerful second-order belief about why the belief would be good to have). But being caused to believe by such a mechanism is not believing at will.⁴⁴ Thus it seems any exercise of control over a belief in response to extrinsic reasons will require a mediating intention.

Of course, the requirement that an intention mediate between extrinsic reasons and a resulting belief hardly threatens the possibility of believing at will. After all, anything done at will requires a “mediating” intention. To show that one can't believe at will, I

must show that one can't form, in response to extrinsic reasons, an "intention to believe" and, in executing that intention, believe as an ordinary action.⁴⁵

We have seen that both ordinary mortals and Credamites can find convincing the extrinsic reasons for believing *p*, therein form an intention regarding the belief that *p*, execute that intention, and end up believing *p*. I will now argue that in executing any such intention—no matter how quickly or efficiently—one will always *act upon oneself*, managerially, to *bring oneself* to believe. Thus, in executing such an intention one will not believe in the way one ordinarily acts—one will not believe at will. Regardless of one's contingent psychological constitution, believing is not an action in its own right, but rather the possible *product* of a distinct managerial action

Consider, again, ordinary action. In performing an ordinary intentional action, one answers for oneself the question of whether to $\square$, therein intends to $\square$, and thus—providing all goes well— $\square$ s intentionally. Even though the intention and the action are distinct—one might intend but fail to act (perhaps due to sudden paralysis)—both the intention and the action are answerable to the same set of reasons, viz., those that bear on whether to $\square$. You are answerable to these reasons just by intending to $\square$ —regardless of whether you succeed in $\square$ -ing. And if $\square$ -ing is a non-basic action, such as preparing dinner or getting a gallon of milk, then each bit of the process is answerable to the constitutive reasons for intending the larger action. Buying a gallon of milk may require going to the store, standing in line, and handing over some cash. Insofar as each of these is part of getting milk, the why-question to which they are vulnerable will be answered, in part, by the reasons bearing on whether to get milk.

Ordinary intentional actions display a uniformity of answerability from intention, through process, to successful execution.

Turning, then, to believing for extrinsic reasons. As an ordinary mortal, you might take certain reasons to show a belief good to have, therein form a managerial intention to *bring yourself* to believe, execute that intention through some process (perhaps by conducting an investigation or taking a pill), and, if you are successful, thereby come to believe. In performing this managerial action, you are answerable, at each stage in the process, to the reasons which bear on whether to bring yourself to believe—to the constitutive reasons for the intention to bring yourself to believe (which are extrinsic reasons for believing). Bringing yourself to believe is thus an ordinary, non-basic, intentional action one might perform in response to extrinsic reasons for believing.⁴⁶

But notice, further, that *if successful* this managerial action—bringing yourself to believe *p*—will *create* answerability to reasons which bear on whether *p*. After all, if successful, this action will create a belief, and in believing *p* one is answerable to such reasons. Importantly, however, neither the managerial intention to bring yourself to believe nor the activities involved in bringing yourself to believe are *themselves* answerable to reasons which bear on whether *p*. You have not yet answered the question whether *p*, and so are not yet answerable to reasons bearing on that question. The managerial intention and action are answerable only to reasons bearing on whether to *bring yourself* believe.⁴⁷ Only if your action is successful will you become answerable to the constitutive reasons for believing.

Thus we now see that believing stands distinct from the managerial activity of bringing oneself to believe, not because the believing stands at the end of a multi-step

process, but rather because bringing yourself to believe and believing involve very different kinds of answerability. Moreover, bringing yourself to believe and believing would continue to involve different kinds of answerability, even if we could eliminate the intermediate process. Even if one could bring oneself to believe just by concluding that it would be good to do so—even if one could bring oneself to believe as a basic action—bringing oneself to believe would remain distinct from its product, believing.

Thus we should not understand the Credamites as believing at will—believing in the way they act—but rather as *bringing themselves* to believe, as a basic action. Notice that, even for a Credamite, the intention formed in response to extrinsic reasons for believing is answerable *only* to such reasons. The Credamite does not become answerable to the constitutive reasons for believing *p* unless he executes his intention successfully. If, despite his best intentions, his special capacity to believe in response to extrinsic reasons fails, and he finds himself unable to believe (the analogue of sudden paralysis), the Credamite will then be answerable *only* to the extrinsic reasons for believing *p*. He would not yet be answerable to reasons which bear on whether *p*, because he would not yet have answered for himself the question of whether *p*. Thus, even for a Credamite, the activity performed in response to extrinsic reasons for believing remains distinct from believing, in a way that intending to raise your arm is not distinct from raising it or standing in the checkout line in order to buy milk is not distinct from buying milk. So, speaking precisely, we should say that the Credamites form and (if all goes well) successfully execute an intention to *bring themselves* to believe. Unlike us, they can form and execute such an intention immediately, as a basic action. But possessing this ability—the ability to bring oneself to believe, as a basic action—does not transform

believing into an action in its own right. Rather, belief is still the product of a distinct managerial activity, when that activity is successful. Even the Credamites do not believe at will.

To restate the argument: we learned from Anscombe that which action one intentionally performs—and indeed (we can now add), more broadly, what it is one actively does—can be identified by the question to which one is answerable, in doing it. The question to which you are answerable will be a question you have answered for yourself. By taking certain considerations to be reasons, you have already determined the question on which they bear (as we saw earlier), and so have already determined which question you will answer for yourself in response to those reasons, and so have already determined which activity you will perform in response to them. Extrinsic reasons are, by definition, those you take to bear on whether the belief that *p* is good to have without showing that *p*. Thus, by responding to them you will not answer for yourself the question of whether *p*, and so will not be answerable to that question. But to believe *p* is to be answerable to reasons bearing on whether *p*. So, in responding to extrinsic reasons, you will not believe. You will rather do something answerable to reasons showing the belief good to have—you may intend to make yourself believe, and so engage in the managerial activity of bringing yourself to believe. Given the right psychology, you might even execute this intention as a basic action. But you would not thereby believe at will—you would not thereby believe in the way you act. Rather, you would perform a managerial action productive of belief. You would bring yourself to believe.

In order to believe “immediately” in response to extrinsic reasons—that is, without acting upon yourself managerially—you would have to answer (rather than merely make

yourself answer) the question of whether p for extrinsic reasons. But the extrinsic reasons are precisely those reasons you do *not* take to answer that question. You take them to bear, instead, on the question of whether the belief is good to have, and you will answer *that* question in response to them. Thus believing immediately for extrinsic reasons is impossible. Believing at will, so understood, is thus impossible. The same fact which makes it the case that you must believe “at will” (rather than by evaluative control)—viz., that your reasons for believing are only extrinsic—also guarantees that you cannot.⁴⁸

What of intending at will? Believing, we have seen, cannot be an ordinary action because believing involves a commitment to its content as true. But an intention also involves a commitment—a commitment to act. Thus, in the properly parallel case, you can no more intend at will than believe at will.⁴⁹

There is, of course, a sense in which we *can* intend at will: we can decide to $\square$ and therein intend to $\square$. We can intend to $\square$ by evaluative control, in response to reasons for $\square$ -ing. (In this sense we can also believe at will—we can believe by evaluative control, in response to reasons which bear on whether p .) But the properly parallel case of “intending at will” is not this, but rather intending in response to *extrinsic* reasons—intending immediately in response to reasons which only show the intention good to have, without showing the action to be done. I will now argue that you can no more intend for extrinsic reasons than you can believe for extrinsic reasons. Thus, in the properly parallel case, you can no more intend at will than you can believe at will.

Suppose I have extrinsic reasons for intending to retaliate: intending to retaliate would deter attack, but deterring attack is not, I maintain, a reason to retaliate. I will not

intend to retaliate simply by finding these extrinsic reasons convincing. Rather, by finding such reasons convincing, I may form a *second-order* intention—an intention to (bring myself to) intend to retaliate. Perhaps I execute this intention through some process—perhaps I take a pill, or perhaps I have a Credamitean ability to execute this second-order intention as a basic action. In either case, in forming and executing the second-order intention I will be answerable to the constitutive reasons for that *second-order* intention—reasons which show it good to intend—which are extrinsic reasons for the first-order intention to retaliate. I will become answerable to the reasons that bear on whether to retaliate only if I in fact intend to retaliate—that is, only if I execute my second-order intention successfully. Again, the action performed in executing the second-order intention remains distinct from the intention produced, because it involves very different answerability. Even if, like a Credamite, I could skip the intermediate process, I will not intend in the way I act. Rather, I will act upon myself to make myself intend.

In order to intend to $\square$ in response to extrinsic reasons *without* acting upon yourself to make yourself intend, you would have to take the extrinsic reasons to answer the question of whether to $\square$. But extrinsic reasons are precisely those you think show the intention to $\square$ good to have *without* bearing on whether to $\square$. One cannot answer a question for reasons that one does not, oneself, take to bear on it. One can at most make oneself answer that question. So once again, the same fact which requires that I form a second-order intention, and so qualifies this as a case of “intending at will”—namely, the fact that my reasons for intending to $\square$ are merely extrinsic—also guarantees that the first-order intention that I may successfully form will remain distinct from, and merely a

product of, my executing the second-order intention. That is to say, the same fact that qualifies this as a case in which I must “intend at will” (rather than by evaluative control) also guarantees that I will be acting upon myself, managerially, and so only making myself intend.

CONCLUSION

I have tried to show why believing is not voluntary—why one cannot believe in response to practical reasons in the way one acts in response to practical reasons. This inability is, as suspected by Williams and Bennett, a conceptual matter, traceable to the fact that belief aims at truth—in particular, to the fact that to believe p is to be committed to p as true. Thus, in believing p , you are answerable to the question of whether p . But the reasons for which you must believe, if you are to qualify as believing at will, are precisely those that you do not take to answer this question. Thus you cannot, for these reasons, answer that question and incur the commitment. You can, at most, bring yourself to answer the question and incur the commitment. That is, you can, at most, bring yourself to believe. In performing an ordinary action, there is no need to act upon yourself in this way—there is no need to bring yourself to answer a question other than the one on which one’s reasons for acting bear. Thus, although you can bring yourself to believe for practical reasons, you cannot believe in the way you act. You cannot believe at will.

I also tried to show that intending involves a similar commitment, a commitment to act. Therefore, you can no more intend for reasons that only show the intention in some way good to have than you can believe for reasons that only show the belief in some way good to have. You can at most bring yourself to intend. Thus, even though belief and

intention are subject to two quite robust forms of agency—evaluative and managerial control—one can no more intend at will than believe at will.⁵⁰

Department of Philosophy
University of California, Los Angeles

¹ Bernard Williams, “Deciding to Believe,” *Problems of the Self* (Cambridge: Cambridge University Press, 1973): 148. Originally published in H. Kiefer and M. Munitz, eds., *Language, Belief, and Metaphysics* (Albany: State University of New York Press, 1970).

² In fact, Williams argument makes use of a premise about needing *knowledge* of one’s ability, a detail picked up by Bennett. As this detail is unimportant for my discussion, I have suppressed it.

³ Jonathan Bennett, “Why Is Belief Involuntary?,” *Analysis* 50 (1990): 93. The resources for this response were put forward by Barbara Winters, in “Willing to Believe,” *Journal of Philosophy* 76 (1979): 243–256. A further argument was developed by Dion Scott-Kakures, in “On Belief and the Captivity of the Will,” *Philosophy and Phenomenological Research* 53.4 (1993): 77–103. Scott-Kakures’ account is criticized by Dana Radcliffe, in “Scott-Kakures on Believing at Will,” *Philosophy and Phenomenological Research* 57.1 (1997): 145–151. Joseph Raz has a helpful discussion in “When We are Ourselves,” *Engaging Reason* (Oxford: Oxford University Press, 1999).

Williams has another argument to the same conclusion, which Bennett also seems to counter successfully. However, considering that argument is not necessary for my purposes.

⁴ It is worth noting that Williams’ claim that “belief aims at truth” has been interestingly developed by both J. David Velleman and Ralph Wedgwood. However, the claim that belief is involuntary *because* it aims at truth does not fare any better on either development. Velleman understands the aim of belief in terms of the functioning of the psychological systems that regulate it: belief aims at truth because it is a state that is regulated to track truth. (See his “On the Aim of Belief” and his “Introduction,” both in his *The Possibility of Practical Reason* [Oxford: Oxford University Press, 2000].) However, it is unclear why the Credamites could not bring themselves to have a state that is thereafter regulated for truth, and so bring themselves to believe. Wedgwood has suggested that the aim of belief is a normative one: a belief is correct only if its content is true. Again, it is unclear why a normative “aim” such as this would prevent the Credamites from turning their trick. (See his “The Aim of Belief” *Philosophical Perspectives* 16 [2002].)

In a recent paper Nishi Shah has sided with Wedgwood, against Velleman, suggesting that the normative aim of truth, together with a necessary disposition to believe what one ought, will explain the “transparency” of first personal deliberation. See his “How Truth Governs Belief,” *The Philosophical Review* (forthcoming). If Shah were right, it would be impossible to believe at will due to this disposition. Having the disposition, Shah suggests, is required for having the concept of belief. This seems to me both too strong—I don’t see why some creature (like a Credamite) couldn’t have the concept of belief (perhaps noticing that most believers have this disposition) without itself possessing the disposition—and therefore too weak—it would still render the inability to believe for practical reasons a merely psychological fact.

⁵ Bennett, 90.

⁶ Bennett treats “practical reasons” as equivalent to what he calls with “inducements.” Of inducements he says, “Quite generally, inducements should be thought of as inducements to *make it the case that p* for various values of *p*” (89).

⁷ It is worth noting the inclusiveness of this claim: an activity is voluntary if it *can* be done immediately for practical reasons. Perhaps we sometimes act without reason, and do so voluntarily. Still, the action could have been done immediately for practical reasons.

⁸ This example comes from Richard Feldman, “The Ethics of Belief,” *Philosophy and Phenomenological Research* 60.3 (2000): 671–2.

⁹ Davidson consider similar cases in which one is caused to do something for a practical reason without yourself doing it intentionally, for those reasons. In Davidson’s example, a climber thinks he has a reason

to loosen his hold on the rope by which he supports the weight of another— because doing so will free him from danger—and this thought so unnerves him that he in fact loosens his hold. But, though he here “responded to reasons,” he did not do so in a way that amounts to responding intentionally, or actively, or “at will.” Donald Davidson, “Freedom to Act” in *Essays on Freedom of Action*, ed. Ted Honderich (New York: Routledge and Kegan Paul, 1973): 153–4.

¹⁰ Bennett wants to show that belief can’t be what he calls “motivationally” immediate to practical reasons, in the way that actions are. He allows that even actions are what he calls “ontologically” mediated: when you respond to practical reasons by, say, immediately raising your hand, there are events that mediate between your appreciation of the reasons and your action (certain neurons firing and muscles contracting). The relationship between your reasons and your action is thus ontologically mediated. However, there is no further thing that you are motivated to do, to bring it about that you raise your hand. So, the relationship between your reasons and your action is motivationally immediate. In contrast, Bennett wants to show that practical reasons cannot immediately motivate belief. They must motivate you to do something else, which in turn results in your believing.

I will suggest that practical reasons are “motivationally immediate” to intending, rather than believing, as a conceptual matter. So practical reasons for a belief will always require some mediating intention. Moreover, the intention they might motivate, I will argue, should not be understood as an “intention to believe,” where believing has become an action, but rather an “intention to make oneself believe.” This is so even if one can execute that intention as a basic action.

¹¹ See J. David Velleman’s “On the Aim of Belief.” My thinking on this subject is indebted both to Velleman’s article and to further helpful correspondence with him.

¹² I do not here mean to suggest that the inability to answer a request for justification will necessarily leave a believer unwarranted, or even unjustified, in her belief.

¹³ Of course, a creature needn’t *itself* have the concept of belief in order to be a believer. Consider small children. Still, if one does have that concept (as one must, if one is to *ascribe* beliefs), then one sees the applicability, in principle, of a certain sort of question. Thanks to Tyler Burge and Gavin Lawrence for help on these matters.

¹⁴ I am not here attempting to offer an analysis of believing. Rather, I am taking for granted our common sense of the standards to which a believer is subject, and inter-defining a number of terms.

Cf. Wedgwood, who suggests that “types of mental state are individuated by the conditions under which they satisfy normative concepts” and that “primitive norms,” such as the norm of truth, “articulate [the] essence or nature of belief. They . . . articulate, as we might put it, constitutive features of belief—that is, features that make belief the type of mental state that it is” (270–1).

¹⁵ I am here partly following T. M. Scanlon, who takes a reason to be a consideration that counts in favor of a judgment-sensitive attitude. See his *What We Owe to Each Other* (Cambridge MA: Harvard University Press, 1998): 17–22. I wholly agree with the thought that a reason is a consideration. I further agree with the thought that a consideration becomes a reason only when it stands (or is taken to stand) in a certain relation. However, I understand a reason, most fundamentally, to be a consideration that bears on a question, rather than one that counts in favor of an attitude, because the latter formulation is ambiguous, in the way I am about to explain. To avoid the ambiguity, one must avoid saying that a reason counts in favor of a judgment-sensitive psychological state, or even, for that matter, a mental activity like believing or intending. Rather, a reason can be thought of as a component in a piece of reasoning. Thus my formulation: a reason is a consideration that bears on a question. See my “Reasons, Actions, and Attitudes” (in progress).

¹⁶ In common usage, “belief” is ambiguous between the attitude (or activity) of believing and the content believed. I will standardly refer to belief as the attitude or activity; when I need to refer to the content, I will make that explicit.

¹⁷ Christian Pillar, “Normative Practical Reasoning,” *Proceedings of the Aristotelian Society*, Supplementary Vol. 75.1 (2001): 195–216. I believe Pillar is making the same distinction introduced to me by Derek Parfit, using the labels “object-given” and “state-given.” See his “Rationality and Reasons” in *Exploring Practical Philosophy*, ed. Dan Egonsson, et al. (Burlington, VT: Ashgate Press, 2001). Neither author says precisely how the distinction is to be drawn.

¹⁸ Though, as we will soon see, by finding the content-related reasons convincing one has already believed, so the fact that the belief is good to have will be, in such a case, inert—a happy fact, but nothing to act upon. (In cases of irrationality, you may have to bring yourself to fully believe something you already think true, because believing the truth fully would be good to do.)

¹⁹ In addition to the peculiar, “This belief is good to have,” consider the following case: Suppose you are interviewing people to run your ad campaign, and you find yourself thinking it would be good to believe that a certain person would be best for the job—you find yourself thinking you’d seem younger and sexier if only you could believe that she was the best person to hire. These considerations show something good about believing she is the person to hire, but such considerations do not usually bear on whether this person is most qualified candidate. Yet, if you find yourself thinking the belief would be good to have as a result of the candidate’s own doing, then the fact that it would be good to believe that she will be an effective leader of your ad campaign might, itself, be a reason that shows that she will. Here, *p* is a proposition whose truth you think is made more likely by the fact that you have a reason for wanting to believe it.

²⁰ I take these to be equivalent and find it noteworthy that the active and passive voice apply equally.

²¹ This is a substantive claim, though I think an obvious one. It is substantive, rather than tautologous, because finding these reasons convincing isn’t just the same thing as believing *p*. Perhaps the Credamites bring themselves to believe *p*—perhaps they bring themselves to take *p* to be true in a way that leaves them answerable to the belief-relevant questions—without finding convincing any reasons that bear on whether *p*. Indeed, I don’t think we need to appeal to the Credamitean ability for this point. Sometimes we find ourselves with beliefs for which we, at least at the moment, lack convincing reasons—we find ourselves taking something to be true in the relevant way. (Perhaps a belief cannot sustain reflection upon the fact that you lack reasons for it. Still, one might unreflectively sustain a belief for which one now lacks convincing reasons, a belief which might inform one’s actions, and that should be enough to show that there is a difference between being committed to *p* as true and finding convincing reasons that bear on whether *p*.)

Some might want to add: “insofar as you are rational.” While this might be an important addition for intention, I don’t think it necessary for belief. If I find the reasons truly convincing, I will believe. If I don’t believe, then I haven’t found them convincing. I might, of course, find myself with conflicting or incoherent sets of beliefs, and so be subject to a charge of irrationality.

²² As noted in this paragraph, constitutive reasons for a belief that *p* are those *taken to* bear on whether *p*. It does not matter whether they “really” bear on whether *p*—so long as the believer takes them so to bear, she will believe by finding them convincing, and they will be what grounds her commitment to *p* as true. I will sometimes take this point for granted, as constantly making note of it generates considerable clutter.

Note, too, that the constitutive reasons are not themselves constitutive of the belief. One can maintain a belief while forgetting one’s reasons for it. They earn the title by supporting the commitment that is constitutive of the belief.

²³ The relation between these distinctions is described, in this paragraph, from a single point of view: those considerations taken to be content-related reasons will also be constitutive reasons, the remaining reasons, the extrinsic ones, will be taken to be attitude-related. People may of course disagree about whether a particular consideration is a content- or only an attitude-related reason, and so disagree about whether that consideration should provide a constitutive or only an extrinsic reason for the belief.

²⁴ One might think that extrinsic reasons are not “really” reasons for the belief, that the only real reasons are the constitutive ones (or, to be precise, the reasons that would be constitutive for someone who was correct about them). While there is no doubt truth to this, I think it unhelpful to restrict the word “reason” to the constitutive ones. I address this in my “Reasons, Actions, and Attitudes.” (For a nice discussion of surrounding issues, see Thomas P. Kelly “Epistemic Rationality as Instrumental Rationality: A Critique” *Philosophy and Phenomenological Research* 66.3 [2003].)

²⁵ This distinction is, I think, a development of points made by Richard Moran. See his *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton: Princeton University Press, 2001): 116–19.

²⁶ I do not think that these two forms of control exhaust of all possible forms of control. More needs to be said about the kind of control we exercise over our own bodies, for example, or over other people. These two forms do seem to me to exhaust the forms of control that we exercise over our beliefs and intentions.

In an interesting recent book, David Owens discusses our control over our beliefs and its relation to our responsibility for our beliefs. I hope to take up issues of responsibility and what he calls “reflective control” in subsequent work. See David Owens, *Reason without Freedom: The problem of epistemic normativity* (London: Routledge, 2000).

²⁷ One might object that whether a person believes *p* is determined, in part, by how that person goes on to act. I agree. But this is only to say that whether that person has answered for herself the question of whether *p* depends on how she goes on to act.

²⁸ The point about indirection, while useful for introducing the distinction, may not be a distinguishing feature of exercises of evaluative control. The immediacy of evaluative control is the deeper feature, as it is a direct consequence of the commitment-constitution of belief.

My claims about the commitment-constitution of belief and the immediacy of evaluative control should not be confused with Moran’s observations about what he calls “transparency” when discussing first-personal knowledge of beliefs (in *Authority and Estrangement*). Rather, what I have to say might be understood as explaining Moran’s observations.

Transparency, for Moran, is a relation between *questions*. The “theoretical” (psychological) question “do I believe *p*?” is transparent to the “deliberative” question “is *p* true?” in that, in the normal case, one answers the first question by answering the second. In fact, Moran claims that part of what it is to be rational is to answer the first question by answering the second—i.e., to meet the “transparency condition” when avowing one’s beliefs. (This, in turn, is to explain how we can know our own beliefs “immediately,” i.e., without evidence about our own psychology.)

The immediacy of evaluative control, in contrast, does not concern the relation between two questions (nor one’s “immediate” knowledge of one’s own beliefs), but rather concerns the relation between answering a question and having an attitude: if you answer for yourself positively the question of whether *p*, you immediately—therein, *ipso facto*—believe *p*. Immediacy, as a relation between answering a question and having an attitude, can explain how transparency could be a *rational* relation between a theoretical and a deliberative question: the immediacy of evaluative control ensures that answering the second, deliberative question, “is *p* true?” will *make true* an answer to the first, theoretical question “do I believe *p*?”

The immediacy of evaluative control, in turn, must be explained by a fact about the attitude itself. How is it that, by answering a question about the world you *make true* the answer to a question about yourself—that is, how is it that, by answering the question about whether *p* you *ipso facto* believe *p*? This is so because to believe *p* is to be committed to *p* as true—to take *p* to be true in such a way as to be answerable to certain questions and criticisms. By answering the question of whether *p*, you take *p* to be true in precisely the way that leaves you thus answerable, i.e., you believe.

So I take Moran’s observations to reveal something about the nature of belief. How could it be true that we can rationally answer a question about our psychology by answering a question about the world? This would be so if, by answering the second question, we make true an answer to the first question. Answering the second question would make true an answer to the first if belief is commitment-constituted, because then, by answering the second question, one would, therein, believe.

My explanation of transparency thus differs sharply from that offered by Shah in “How Truth Governs Belief.” Shah hopes to explain a transparency quite different from that which concerns Moran. Shah hopes to explain why, “in asking oneself whether to believe *p*, one is forced to recognize that this question is answered by settling the question of *p*’s truth” (16). His answer is that “possessing the concept of belief involves being disposed [to be moved only by truth-sensitive considerations] when one applies the concept to one’s own thinking” (34). I think the invocation of this motivational disposition is quite unnecessary to explain the phenomena at which Shah is pointing. If we allow that competent English speakers understand that to say a person believes something is to say that that person thinks that thing true, the phenomena has a rather boring, but otherwise entirely satisfactory answer: as a speaker of English, you know that, in your own case, to ask whether to believe *p* just is to ask whether to think *p* is true, and so considerations that bear on whether *p* is true have obvious relevance. Shah makes the overstrong assumption that, in asking whether to believe *p*, one is *forced* to appeal to considerations that bear on whether *p*, and he invokes his disposition to explain this supposed fact. But, to my mind, he is explaining something that isn’t so. In asking whether to believe *p*, I am not forced to appeal to considerations that bear on whether *p*. I may take up that question in an altogether different spirit, asking whether it would be useful to believe *p*. Of course,

as I am about to explain, it will turn out that I will not believe simply as a result of answering this question. But this is not due to any special motivational disposition that comes with the concept of belief, but rather due to the more mundane conceptual fact that to believe just is to think true. (Perhaps this mundane fact creates “motivational dispositions” in those competent with the concept—but I think it misleading to make the disposition do the work.)

²⁹ I owe this helpful example to Tom Kelly, who uses it in his discussion of the rationality of belief. See note 26.

It is worth noting the somewhat odd character of this investigation. Not all investigations are prompted by extrinsic reasons. Simple concern for your friends, for example, might prompt the call. However, this doesn’t seem to be grounded in the desirability of having a certain belief. I’m not sure what to say about the relation between concern and placing the call. Alternatively, one might object to the example by saying, “I don’t just want to believe my friends are home, I want to *know* it!” I take this to mean that, you don’t just want to believe that they are home, you *also* want it to be true that they are home. Still, a part of what you want is to have the belief, perhaps on the condition of its truth. And that should be enough to make the point at hand. For an alternative example, suppose you are anxious about the strength of the housing market, even though you have already decided to make a purchase. You might do some research, just to quell your anxiety.

³⁰ While one can, of course, exercise self-management aimed at arriving at *some* true belief about *whether* *p*, I mean for the management I am discussing to aim at a *particular* belief, that *p*. One might object, thinking that you can only intend to do things you think you are able to do, and so, if you don’t know whether there is good evidence that *p*, then you can’t intend to believe *p*, in particular, but only to arrive at some belief about whether or not *p*. But if I can intend to shoot a free-throw, given what I know of my chances (and not just to *try* to shoot a free-throw), then it seems that I can, at least in many cases, intend to bring myself to believe that *p*. At any rate, once we introduce the science-fictional possibilities, our ability to intend to bring ourselves to a particular belief should be clear.

³¹ The term “manipulative control” is more apt for those cases in which you bring yourself to a belief *without* providing yourself with good constitutive reasons, while “managerial control” is more apt for those cases in which you bring yourself to belief by providing for yourself good constitutive reasons. For the purposes of this paper, however, we need not distinguish between these cases—however you do it, you are exercising manipulative/managerial, as opposed to evaluative, control.

³² Nor do I claim that it is possible. I do not claim that imagination is a reliable guide to possibility. Perhaps a thing cannot be a belief if it is brought about in these ways, or perhaps such a process could not secure the requisite form of answerability. I have nothing to say about these questions.

³³ Cf. Gilbert Harman, *Change in View* (Cambridge: MIT Press, 1986): 32–35.

³⁴ The caveats from earlier apply here as well. One cannot actually be open to requests and questions without being a self-reflective, language-using creature. However, the applicability of the concept depends on the questions being, in principle, in place. This is because having the attitude leaves one vulnerable to criticism under certain standards.

³⁵ For an extended discussion of intention, including some discussion of the commitment involved in intending, see Michael Bratman’s *Intention, Plans, and Practical Reason* (Cambridge: Harvard University Press, 1987). Some of his subsequent development can be found in his *Faces of Intention* (Cambridge: Cambridge University Press, 1999).

³⁶ Cf. G. E. M. Anscombe, *Intention* (Oxford: Basil Blackwell, 1957). As we will see, Anscombe notes that any intentional action leaves one open to a certain sort of “why question.”

Of course, one way that I can do something intentionally is by intending to do some larger action of which it is a part. In this case, I will say that the intention to do the larger action included an intention to do the part. (Sometimes I unintentionally do something—in ignorance, say—in the course of doing something else intentionally. In such a case, I intend the larger action, but not the piece I do unintentionally.)

³⁷ In saying one “answer[s] for oneself (however implicitly) the question of whether to $\square$ ” I do not mean to imply that we should be able to find, by psychological experiment, some underlying process that corresponds to this answering of a question. Rather, we can posit the “implicit” answering on the basis of

the answerability. If one intends, one is answerable in just the way one would be if one had explicitly answered the question of whether to $\square$

³⁸ Pleasing my boss both bears on whether to do the job and, I am supposing, is reason enough to do it. In other cases, reasons for having the intention will bear on whether to act but will not be reason enough to act. If, in the barroom, you offer me five dollars to intend to down some whiskey, I can decide to down the whiskey, thus have the intention and win the money. But if you offer me five dollars to jump from the second story window, I cannot then decide to jump, and so cannot win. Although the offered money is certainly reason enough to (merely) intend to jump, and though the offered money seems also to bear on the question of whether to jump, it is not reason enough to jump. The overall reasons against jumping far outweigh the reasons for jumping.

These cases contrast with Gregory Kavka's toxin puzzle ("The Toxin Puzzle," *Analysis* 43 [1983]: 33–36). In the toxin puzzle, you are offered a very large sum of money to intend, at midnight, to drink a non-lethal but temporarily very unpleasant toxin tomorrow afternoon. If you intend tonight, he will pay you the money tomorrow morning, and you will be permitted to keep the money regardless of whether you follow through on your intention and drink. You are being paid to intend, not to actually drink it. Further, you are prohibited from taking steps to give yourself other reasons for drinking tomorrow (such as establishing side-bets with friends about whether you will drink), and it is stipulated that you are not "one of those strange people who take pride in never releasing oneself from a promise to oneself, no matter what the circumstances" (35).

The toxin case is like the case of downing whiskey and unlike the case of jumping out the window in that the reasons for having an intention are much stronger than the reasons against acting. Yet it seems that, in the toxin case, you cannot decide to act in order to have the intention. I think the puzzle has this form because the strong reasons for having the intention do not even *bear on* the question of whether to act. Kavka carefully builds his toxin puzzle so that, when the time comes to drink, you will no longer have any reason even to *have the intention* to drink, and you know this now. Thus, in Kavka's puzzle, the fact that you have reasons showing the intention useful does not provide you with reason to act—and so it seems not even to bear on the question of whether to act—due to the delay between when the intention is useful and when the action is to be performed. The fact that the reasons for intending are not reasons for acting explains why, even though the rewards of intending to drink are much greater than the penalty for drinking, you cannot decide to drink.

This interpretation of the toxin case explains the inability to decide to drink the toxin by appeal to a lack of reasons for drinking, and explains that lack by appeal to the delay between the time at which the intention is useful and the time at which the action is to be performed. If this interpretation is correct, it should apply equally well to what John Broome has called the "non-toxin case." In this alternative case, you are offered money to intend to drink a glass of water at noon tomorrow. It might seem that you can decide to drink the water tomorrow, and so win the money. But on the interpretation that I have provided, you should have no reason to drink at noon: after all, by noon, you will no longer have reason for intending to drink the water, and you know this now. On my interpretation, whether you can win the money in the non-toxin case would have to turn on whether you can decide, now, to do something, tomorrow, that you will have no reason, tomorrow, to do.

I allow that one can decide to perform certain actions with very little justification—so long as there are no reasons against doing them, they can be performed on whim, or just because it occurs to you. Drinking a glass of water is presumably among these. Perhaps, then, you can decide, now, to drink tomorrow, simply to suit your fancy, and so win the money.

To be honest, though, my confidence that I can decide to drink the water tomorrow seems subject to the same slow erosion by reflection that undermined my initial confidence that I could decide to drink the toxin tomorrow. I will win the money now only if I truly intend to drink tomorrow. Suppose I decide, now, to drink tomorrow, just to suit my fancy. I cannot appeal to reasons of consistency or loyalty or guilt to ensure that I will follow through on this decision—such reasons are ruled out of the original toxin case. So I should know, now, that the slightest reason against drinking, or even a change in the winds of my fancy, will prevent me from drinking tomorrow. How strong, then, must the intention be, for me to win? If a very weak intention will do, perhaps I can succeed. But I am not sure that I can decide, tonight, to do anything

on a whim tomorrow—I don’t know that whims can be so enduring. (Perhaps my intuitions about this case simply betray an insufficiently robust sense of whimsy.)

In any case, at this point in the text, I am concerned to account for the fact that sometimes it seems you can decide to act because you have reason to intend to act. It is enough, for this point, if the fact that I have reason for the intention sometimes provides me with reason to act. Many thanks to John Broome for raising this interesting issue.

³⁹ This account of intentional action is, of course, just a skeleton. I have simplified in many ways and am overlooking many complications. However, for the purposes at hand, I believe it will serve.

With this account in place, we can characterize acting *immediately* in response to practical reasons: to act immediately in response to practical reasons would be to form, in response to such reasons, by an exercise of evaluative control, an intention to do a basic action right now, and then, right now, to execute that intention and perform the action. But it is unclear why immediacy of this sort should render an action more or less voluntary.

⁴⁰ *Intention*, 9. We could just as readily say that, in acting, one is vulnerable to criticism or critique under certain standards, and that, in reflective, language using creatures, this leaves us open to certain questions. Of course, the standards to which one is vulnerable, in acting, are less clear than those to which one is vulnerable, in believing. Are the standards set by one’s own ends and beliefs? Do moral standards always apply? We needn’t answer these questions here. It is enough that in acting intentionally, as in believing, I commit myself in a way that leaves me answerable to a certain sort of why-question.

⁴¹ One is “in a position to answer” a question, here, even if one’s answer is poor—a poor answer is still an answer; likewise, as Anscombe notes, the answer “for no particular reason,” it is still an answer to, rather than a rejection of, the relevant why-question (9).

My account of why one is in a position to answer Anscombe’s why-question, as well as the account I would give of Anscombe’s second mark of intention action—that one knows what it is that one is intentionally doing—thus differs from the account offered by Kieran Setiya, in his “Explaining Action,” *The Philosophical Review* (forthcoming). Setiya starts his paper with a terrific question: why should Anscombe’s two marks of intentional action show up together? As with my earlier difference with Shah, my own answer to this question seems to me more boring, but otherwise entirely satisfactory: if to intend just is to have answered positively the question of whether to $\square$, then Anscombe’s why question is “given application” whenever one intends just because that question looks for just those reasons which would answer the question you have already answered, in intending. Likewise, since the question one has answered, in intending, makes explicit mention of what one intends to do, one should be ready to say what it is that one is doing, whenever one acts intentionally (so long as one knows of one’s intention—unlike Setiya, I would allow unconscious or self-deceived intentions). There is no need to recourse to Setiya’s sophisticated solution, which packs an explanation of the motivational basis of one’s action into the content of one’s (self-referential) intention. (Indeed, Setiya’s explanation would, I think, grant us far too much knowledge—for any intentional action, we would know not only what we are doing, but also why. I do not think that motivational clarity is part of the very concept of intentional action. I will not here engage with Setiya’s complex argument for his view, but only note that I do not believe his own view is required to answer the terrific question with which he starts his paper.)

⁴² This is the sense in which we can intend “at will”: we can intend in response to practical reasons without having to act upon ourselves to make ourselves intend. We can find convincing reasons that bear on whether to $\square$ and therein, *ipso facto*, intend to $\square$. We “intend at will” by exercising evaluative control.

⁴³ Though this argument has not, to my knowledge, been previously advanced, the point has been acknowledged by some. Cf., e.g., John Broome: “Reasoning is validated by its contents, and the mental states have to go along willy-nilly” (“Normative Practical Reasoning,” *Proceedings of the Aristotelian Society* Supplementary Volume 75:1 [2001]: 186). Joseph Raz makes a related point about action: “I cannot choose to have coffee because I love Sophocles” (“When We are Ourselves” *Engaging Reason* [Oxford: Oxford University Press, 1999]: 8).

⁴⁴ I take it as a premise that, if one does not believe by evaluative control nor intentionally, one did not believe by an exercise of control of any kind. This seems to me both true and interesting: if you find convincing the reasons which you take to bear on whether *p*, and, thus, straightaway (therein), believe, you

have, yourself, believed—by evaluative control. But if you find convincing other reasons and, straightaway, believe, you have simply been caused to believe by some psychological mechanism.

⁴⁵ This point was helpfully pressed by Mark Greenberg.

⁴⁶ Blushing could easily be an ordinary intentional action for a creature differently constituted. Were we rightly wired, we might sometimes take ourselves to have reason enough to blush, intend to blush, blush, and so be answerable to the question, “why did you blush?” (or, if one happened to be unsuccessful, answerable to the question “why did you intend to blush?”) where that question is looking for constitutive reasons for the intention to blush.

⁴⁷ I don’t mean to deny that this question might need to take into account the likelihood of *p*, such that reasons which bear on whether *p* also bear on whether to bring yourself to believe. In fact, one could accept my point while insisting that a person is *never* justified in bringing herself to believe something false (and so self-management would be appropriate only in cases of irrationality—only when you find yourself not fully believing something you think is true). My point is simply that the truth of *p* bears on the managerial intention and action *only insofar as* it bears on the question of whether to make yourself believe.

⁴⁸ The *constitutive* reasons for believing *p*—those that bear on whether *p*—might be taken to bear on the question of whether to believe *p*. However, one would not, for constitutive reasons, form an *intention* to believe. One would simply believe straightaway, by evaluative control. The only reasons for which one would form an *intention* to believe *p* would be extrinsic reasons—reasons showing something to be said for believing *p* without showing *p* to be true. But for these one can at most make oneself believe.

⁴⁹ In fact, I think the same distinction in agency and limitation of “willing” can be established for any “commitment-constituted attitude”—any attitude for which there are both constitutive and extrinsic reasons. I would argue, e.g., that you cannot resent or forgive at will, either. (See my “Articulating an Uncompromising Forgiveness” *Philosophy and Phenomenological Research* 62.3 [May 2001]: 529–55 and “The Force and Fairness of Blame” *Philosophical Perspectives* 18, *Ethics* (2004): [in press].) Further, a similar point can be made for virtuous or vicious intentions. You cannot form a kind or generous or spiteful intention for reasons extrinsic to those virtues or vices. See my “Extrinsic Reasons, Alienation, and Moral Philosophy” (in progress).

⁵⁰ I am grateful for the help of many. I received on-going input from Tyler Burge, Barbara Herman, Thomas P. Kelly, Sean Kelsey, Gavin Lawrence, Richard Moran, Seana Valentine Shiffrin, and Gary Watson. I had helpful conversation and/or written comments from many others, including Jonathan Adler, John Broome, Luca Ferrero, Harry Frankfurt, Mark Greenberg, Gilbert Harman, Paul Hoffman, Andrew Hsu, Aaron James, Mark Johnson, David Kaplan, Aimée Koeplin, Matthew Lockard, Brad McHose, David Merli, Elijah Millgram, Anselm Müller, Calvin Normore, Nikolaj Nottelmann, Thi Nyguen, Charles Otwell, Derek Parfit, Terrence Parsons, Andrew Reisner, Abraham Roth, T. M. Scanlon, Nishi Shah, Nicholas Silins, Rebecca Slayton, Sheldon Smith, Christopher Smeenk, Frederick Stoutland, Julie Tannenbaum, J. David Velleman, R. Jay Wallace, anonymous reviewers, members of the Southern California Law and Philosophy Discussion Group, my graduate seminar at UCLA in 2001, the autonomy workshop at UC Riverside, an audience at the Princeton Philosophy Department, and an audience at the Southern California Philosophy Conference. Work on this paper was generously supported by the NEH and the American Council of Learned Societies. Any views, findings, conclusions, or recommendations expressed in this paper do not necessarily reflect those of the National Endowment for the Humanities or the ACLS.