
HUMEAN HEROISM: VALUE COMMITMENTS AND THE SOURCE OF NORMATIVITY

BY

VALERIE TIBERIUS

Abstract: This paper addresses the question “In virtue of what do practical reasons have normative force or justificatory power?” There seems to be good reason to doubt that desires are the source of normativity. However, I argue that the reasons to be suspicious of desire-based accounts of normativity can be overcome by a sufficiently sophisticated account. The position I defend in this paper is one according to which desires, or more generally, pro-attitudes, do constitute values and provide rational justifications of actions when they are organized in the right way.

1. *Introduction*

Skepticism about practical reason used to be confined, for the most part, to a concern about the legitimacy of moral reasons. This skepticism is the result of the thought that moral reasons are not obviously connected to individual motivational states, in combination with the popular position that reasons are the kinds of things for which people can act, and therefore the kind of thing that require the existence of a motivational state.¹ More recently, the legitimacy of hypothetical or instrumental reasons has been the object of skepticism.² For example, in a recent paper, “The Normativity of Instrumental Reason,” Korsgaard argues persuasively that the hypothetical imperative cannot be the only norm of practical reasoning because it does not provide us with a reason to pursue our ends in the first place.³ Unless we have some reason to pursue our ends, there seems to be no reason to take the means to them, and if there is no reason to take the means to our ends, then the instrumental norm has no reason-giving or normative force.

If this argument is correct, the instrumental norm cannot be the “source of normativity”; that is, it cannot generate justifying reasons on its own. The normative force or authority of practical reasons has to come from somewhere else, and Korsgaard thinks that there are two alternatives. If some reasons have normative force, then either they derive it from an unconditioned principle of practical reason, or they have it in virtue of our freely (or authentically) bestowing it on them. Korsgaard calls the second alternative “heroic existentialism.” If we cannot make the case that reasons get normative force in either of these ways, then we are stuck with the view that no reasons ever have normative force; in other words, we are stuck with normative nihilism.⁴

Many philosophers are skeptical about the possibility of defending the Kantian project either because they suspect that the only defensible principles of pure practical reason are too formal to give substantive advice, or because they do not believe that there is any defense of even highly formal principles. I do not intend to argue that Kantians cannot succeed in defending universal principles of practical reason, but I think that the prevailing doubt about this is enough to warrant a full exploration of the alternatives.

Heroic existentialism, although Korsgaard does not describe it in much detail, seems to be the view that the normative force of reasons comes from our simply taking our commitments to be normative for us. This is a kind of heroism, I take it, because it requires a steadfast commitment in the face of the knowledge that there is nothing which requires this commitment, no pure principle of practical reason which makes one commitment any better than another. Heroic existentialism seems to mean creating reasons *ex nihilo* and at the same time maintaining that there are reasons (thus avoiding nihilism). This puts the existential hero on shaky ground, and the ability to stand on this shaky ground might seem heroic. Now this position does not look terribly attractive because it cuts the normative loose from processes of justification and reason giving. According to heroic existentialism, it is enough to say “I choose it”; no further justification of our ends is required, or even available.⁵ This seems to leave us with the unhappy view that when we are looking for reasons, anything goes.

A third alternative, which Korsgaard has already dismissed, is that our desires are normative in themselves. In other words, we have a normative or justifying reason to pursue what we desire just in virtue of our desiring it. Thought of one way, this view seems compelling. After all, when we are asked to explain our reasons for acting, desires are the key component of such explanations. However, when we insist on seeking *justifying* reasons rather than merely explanatory ones, this position looks less plausible. This is because it just does not seem to be the case that any desire can explain why it was good to act as we did; like the existentialist position, this position seems to divorce normativity from justification.

Some of the desires we have are fleeting, some are the result of addiction, some are ones of which we are deeply ashamed and some are simply ones we think we would be better off without. Although these desires may *cause* us to act, we do not take them to give us good, justifying reasons for action.⁶

So, it seems that Korsgaard has good reason to doubt that desires are the source of normativity. Desires are not the same thing as values, and explanations of actions by desires are not the same as rational justifications of actions. However, we might think that the reasons to be suspicious of desire-based accounts of normativity can be overcome by a sufficiently sophisticated account. I am inclined to think this, and the position I defend in this paper is one according to which desires, or more generally, pro-attitudes, do constitute values and provide rational justifications of actions when they are organized in the right way.⁷

The position I defend holds that we can understand the normative force of justifying reasons in terms of our attitudes, commitments, and choices. It is, therefore, in the spirit of heroic existentialism insofar as it neither postulates principles of pure practical reason nor retreats into nihilism, and it gives an important role to choice. Contrary to true existentialism, however, choice is not enough: to have normative force, these attitudes and choices must form a pattern that constitutes our taking ourselves to have reasons for these attitudes and choices. In this paper I will explain what is distinctive about the attitudes, commitments, and choices that we experience as normative, and I will argue that commitments to values that we have for reasons, properly conceived, can play the role of a stopping point for normative inquiry without presupposing unconditional principles of practical reason.⁸

2. *Features of our value commitments: some examples*

At first sight, it seems obvious that our own desires, projects, and commitments give us reasons. I desire to write an article about the normative force of reasons and so I have a reason to sit at the computer and type. I want to avoid being tired tomorrow and so I have a reason to get enough sleep tonight. The problem is that if we mean “reason” in the normative sense, and we are not talking merely about explanatory reasons, then it is plausible to think that not all of our desires, projects, and commitments give us reasons. Some are ones we wish we did not have, and would be better off without. For example, the addict’s desire for heroin, or the fleeting urge to jump from a tall building when one realizes that one could, do not seem like desires that provide reasons.⁹ What we are motivated to pursue does not automatically give us reasons. In Korsgaard’s terms, we do not take every motivational state to be normative for us.

But we certainly have some desires, projects, and commitments that we do take to be normative for us. I will call these desires, projects, and commitments (the ones we take to have normative force) our “value commitments,” or just our “commitments,” for short. I suggest that when we understand the difference between our commitments and the desires and projects that we do not take to be normative, we will be on the road to discovering what gives some of our motivational states normative force. I will not differentiate, in this section, between different types of values a person might have because the account of value commitment I propose is meant to be general enough to characterize a person’s commitment to any kind of value. Some of the examples will have moral significance, others will not. Of course, there are important differences between these different types of value and our commitments to them, and I will discuss some of these differences in section 5. For the purpose of developing a general account of our value commitments we can ignore these complications for now.

There are several distinctive features of the desires, projects and goals we take to be normative. To make this case I want to start with three paradigm examples of characters who have some desire, project, or goal that is normative for them. That is, I want to start by describing some examples that will illustrate what it is to have a value commitment.

First, consider Antonia, the main character from the film *Antonia’s Line*. She has just returned to the village after a long absence, in order to run her family’s farm with her young daughter and no husband. Antonia frequently takes stands against the other farmers and does things that make her unpopular in the small community. She is regarded as an out-cast, and she embraces this reputation. When a very nice farmer with whom she has developed a fine friendship asks her to marry him she refuses because, although she loves this man, she sees marriage as a threat to her independence. Their relationship flourishes to the end of her life, but she never consents to marry him or to live in his house. Antonia not only *is* independent, she also takes pride in her independence; it is normative for her.

For a second example, take Jean Valjean, the main character from Victor Hugo’s *Les Misérables*, who is committed to being merciful. Valjean, an ex-convict, is prevented from returning to crime by a Monsignor whom he met on his first night after being released from prison. The Monsignor shows Valjean mercy and tells him that he is giving Valjean’s soul “back to God.” Valjean promises that he will become a new man, which he does: he changes his identity and becomes a successful businessman and leader of his community. He also becomes a man who is devoted to God, dedicated to keeping his promise to the Monsignor, and firmly committed to being merciful and beneficent to other men and women in the way the Monsignor was to him. When he discovers that a mentally ill man

is accused of being Jean Valjean (who skipped his parole in order to start over) he cannot let this other man suffer unjustly and so he reclaims his identity. When he has the opportunity to kill Javert, the miserable police officer who has devoted his life to hunting Valjean down, Valjean spares him.

Finally, consider Silas Marner, from George Eliot's novel of the same name. At the beginning of the novel Silas has an impoverished set of commitments. He cares for nothing more than the gold and silver coins hidden under his floor, as we can see from the following passage:

He loved the guineas best, but he would not change the silver – the crowns and half-crowns that were his own earnings, begotten by his labour; he loved them all. He spread them all out in heaps and bathed his hands in them; then he counted them and set them up in regular piles, and felt their rounded outline between his thumb and fingers, and thought fondly of the guineas that were only half-earned by his work in his loom, as if they had been unborn children . . .¹⁰

Later in the novel, a small child whose mother has died turns up on his doorstep, and Silas comes to have a commitment to this little girl. Very quickly he feels affection for her, in part because of the sympathy and fellow-feeling generated by his thought that they are both alone in the world and have been brought together for a purpose. When the townspeople, assuming that old Silas would not want to keep the girl, suggest that she should be taken from him he reacts passionately:

“No–no–I can't part with it, I can't let it go,” said Silas, abruptly. “It's come to me – I've a right to keep it.” The proposition to take the child from him had come to Silas quite unexpectedly, and his speech, uttered under a strong sudden impulse, was almost like a revelation to himself: a minute before he had no distinct intention about the child.¹¹

Silas grows fonder and fonder of the girl whom he names “Eppie,” and it becomes clear that his commitment to her is deeper, more sustaining, and more conducive to his happiness than his commitment to his coins ever was.

The first thing to notice is that in paradigm cases of having value commitments we are motivated to pursue or promote that to which we are committed. We generally have pro-attitudes, such as desires and emotions, regarding our commitments. Antonia cares about her independence so much that she is willing to act in such a way as to give up other things she wants for the sake of it. Valjean is willing to risk his life for the sake of living up to his values. Silas cares deeply for Eppie and is moved upon meeting her to prevent others from taking her away.

But there is more to the commitments of these three characters than the motivation to pursue or promote that to which they are committed. If

value commitments were simple motivational or affective states, they could not play the role in our lives that they do. First, our value commitments, in a sense, make us who we are. We identify ourselves in terms of them. Antonia thinks of herself as an independent person, Valjean as a servant of God, Silas Marner as the loving father of Eppie. Second, because these commitments are so important to us, they serve as the basis for planning and for our assessments of how well our lives are going. Not every pro-attitude plays an important role in the planning and assessment of our lives; this is, in part, what it means to say that we do not take all of our pro-attitudes to be normative for us. True value commitments have two other features which allow them to play this role.

First, we need stability in our commitments to compensate for fluctuations in our pro-attitudes. We do feel strongly about our value commitments, but our motivations are shifting and somewhat unstable. Pursuing or promoting our commitments is not always our current or strongest motivation. Sometimes we lose the motivation to pursue these commitments because we are depressed or frustrated, and sometimes because we are distracted by other things. We counteract the influence of depression, frustration, distraction, and other such factors by intending to continue to be motivated in the way we had been motivated without these influences. Further, stable commitments provide parameters within which deliberation typically takes place. Without stable attitudes toward some of our commitments, deliberation would be much more difficult because in every deliberative context everything would be open for consideration; there would be no fixed points to give structure to our reasoning. The lack of stability and structure would make it very difficult to succeed in long term planning for these value commitments.

Diachronic stability, then, is an important feature of the paradigm case of a value commitment. Notice that the commitments in the three examples involve stability across time. All three regard their commitments as providing reasons for future actions; these commitments require planning, and constrain the kinds of plans that are made. If we attend to examples of value commitments we can see that stable value commitments have a presumption in their favor when it comes time for deliberation. Antonia does not deliberate about the value of independence every time she is confronted with any choice that might require her to sacrifice it. The value of independence, for Antonia, is a relatively fixed point in her deliberations, which is given some weight as a matter of course, without her having to explicitly assign it that weight.¹²

Stable value commitments constrain deliberation about future action in such a way that this value is preserved. In other words, we do not deliberate anew about these stable commitments every time: we *intend* to give this commitment weight in future deliberations. Having this intention means that we do not re-evaluate the commitment each time,

and that which options we consider is determined, in part, by the fact that we take the value of this commitment to be given.¹³

The second difference between true value commitments and mere pro-attitudes is that we take the former to be justified in some way. So, a third important feature of these commitments is that one takes oneself to have reasons for them. Granted, we might not always be able to articulate our reasons very well, and we may not have these reasons in mind when we are pursuing our commitments.¹⁴ Still, these commitments are essentially different from whims, fleeting fancies, and stubborn desires in that we take there to be something at stake in these commitments which is not present in other of our affective attitudes.

The role that our value commitments play in our lives makes this sense of justification essential; value commitments could not play the role they do in the planning and assessment of our lives if we did not take there to be reasons for having them. We spend a lot of time and energy on our value commitments, and we sacrifice other things we may want for the sake of them. Valjean risks his own survival for the sake of showing mercy to Javert, and Antonia forgoes the easy comfort of living with the man she loves. It would be very difficult, if not impossible, to make such sacrifices if we did not think that there are reasons to value the things for which we make the sacrifices. We have to be able to conceive of the difficulties in pursuing our ends – the sacrifices and the great expenditures of time and energy – as justified costs, otherwise we would not continue to endure them.¹⁵ As Korsgaard puts it, I must have something to say to myself about why I have the ends I have, “something better, moreover, than the fact that this is what I wanted yesterday.”¹⁶

To understand fully the importance of the last two features of value commitments (diachronic stability and justification) we must understand how they work together. Instability in our value commitments undermines our conviction that we have good reasons for having the commitments we have, and it therefore undermines our sense that any of our commitments ever have normative or guiding force. This is because to take one's commitments to have the authority to guide one's choices and actions one must have the conviction that there is something good about them. This conviction requires confidence in the value of the commitment and, therefore, in the reasons one has for it, as well as a willingness to suspend one's skeptical doubts about it. But it is difficult for this kind of conviction to develop unless there are some stable commitments to support it. Unstable and fluctuating commitments are not generally conducive to developing conviction because we become engaged in doubting the value of the unstable commitments, and this undermines our confidence that any of our commitments provide real normative guidance.¹⁷ Furthermore, having reasons for our commitments increases their stability by providing a justification for continuing to give these commitments

weight in deliberation. There is, therefore, a reciprocal sustaining relationship between the stability in our value commitments and our taking ourselves to have reasons for these commitments.

Paradigm value commitments, then, consist of positive affective states which have diachronic stability, and the conviction that there are reasons for having these stable attitudes. It does seem that people take commitments such as these to be normative, that is, to provide them with justifying reasons for action. My hope is that commitments such as these have the advantage of providing a source of justifying reasons that explains the motivating force of such reasons without allowing that all motivational states provide justifying reasons. Before we can see whether this hope is warranted, we need to make the stable attitude account of value commitments more precise. Moreover, we must see how value commitments, according to the stable attitude account, can generate normative reasons without the help of unconditional principles of practical reason.

3. The stable attitude account of value commitment

Briefly, the account of value commitment I propose says that a paradigm case of a value commitment consists in a pro-attitude, an intention to continue to have this pro-attitude, and the conviction that these attitudes are justified by certain considerations that one takes to be reasons for having that commitment. Let me explain.

In the examples of value commitments discussed earlier, each agent has a pro-attitude that is appropriate to what he or she values. Antonia is fiercely proud of her independence and preserves it by refusing to marry Farmer Bas, Valjean acts mercifully himself and admires mercy in others, and Silas loves Eppie, eventually to the degree that he is willing to give her up if living with her biological father would be best for her. To have a value commitment is, in part, to have certain pro-attitudes with respect to it, but, this is not all it is to have such a commitment. As we saw in the previous section, the values we take to be normative for us are, paradigmatically, stable values that play an important role in our planning, and that determine the shape our lives take. This is why someone who has a fleeting motivation toward something, but no disposition to take this value into account in plans for the future, is not a person we would tend to count as having a value commitment in this normatively significant way. Part of what it is to have this kind of commitment is to intend to continue being so committed.

An intention, following Michael Bratman, is a state whose function is to constrain practical deliberation by establishing standards of relevance for the options that will be considered, and by constraining what counts as an admissible option to choose.¹⁸ A person with the intention to remain

committed will not let momentary fluctuations or lapses in their motivations influence their deliberations about what to do. For example, in deliberating about whether or not to accept the proposal of marriage, Antonia is constrained by valuing her independence. If she were to assess her current motivational states without regard to the attitudes she has had toward her independence over time, we would be inclined to say that she is not really committed to the value of independence.

A value commitment to some end or project, x , requires an intention to continue to be motivated in the relevant way with respect to x . By defining an intention in terms of its function I mean to reject the view that a person must have a conscious intention in the sense of an explicit avowal. Someone whose deliberation is constrained in the right way counts as having an intention, on this view, even if they do not regard themselves as *trying* to remain committed.¹⁹ In other words, the stable attitude account does not imply that we have intentions in the sense of acknowledged goals or plans to include these values in our lives. According to the stable attitude account, one counts as having a value commitment in this sense as long as one's reasoning is in fact constrained. Nor does the requirement of stability preclude rational changes in values: in a deeply reflective moment, one might decide that one no longer has reason to value something one once valued. This is compatible with thinking that when one has a commitment to the value of some end, ordinary instances of practical reasoning are constrained by that value.

The third feature of value commitments is that an agent takes their commitment to be *justified*.²⁰ This is the point at which Korsgaard and others will be inclined to call for Kant. One might think that when we take a commitment to be justified we are assuming that there is some reason for having the commitment that is *independent of one's motivational states and dispositions*. So far it seems that I have only pushed the question about the source of normativity further back. If taking oneself to have reasons requires the justification of pure principles of practical reason, I will not have offered an alternative to the Kantian project. But there is another way of understanding what it is to take oneself to have reasons for having the attitudes one has. Taking an attitude to be justified by reasons is another normative commitment. It is a commitment to the value of a complex pattern of attitudes, namely, the original value commitment remaining stable upon reflection on the considerations in question. Thus in order to understand what it is to take an attitude to be justified by reasons, we can redeploy the analysis of a value commitment.

If we do so, the third component of these commitments is unpacked as follows: What it is for an agent to take it that considerations C justify her attitude A is for it to be the case that she has a pro-attitude toward reflection on C sustaining A and she has an intention, I , to continue to have reflection on C sustain A .²¹ When I say that she has a pro-attitude

toward reflection on C sustaining A, I mean, in other words, that she approves of the following pattern that she finds in her own reflection: thinking about her reasons for valuing x tends to make her feelings about x more stable, rather than undermining them. This is not yet sufficient to capture what it is to take one's commitment to be thoroughly justified, though, because we frequently take ourselves to have *some* reasons for having a commitment even though these reasons are outweighed by other considerations we take to be overriding. So, we must add to the analysis the following clause: it is not the case that there are other considerations, D, which are such that the agent has a pro-attitude toward reflection on these considerations destabilizing the attitudes that constitute the agent's commitment, and she intends to continue to have this pro-attitude toward this pattern of destabilization (in other words, where she regards D as a good reason to abandon her commitment). In short, a person with a value commitment approves of her attitudes remaining stable upon reflection on certain considerations, and is aware of no considerations such that she approve of her attitudes becoming unstable in response to reflection on them.

We can notice at this point that I have not applied the analysis of valuing in its entirety to the state of taking an attitude to be justified. I have not said that to take her attitude to be justified the agent must take the pro-attitude she has toward "A remaining stable upon reflection on C" to be justified by other considerations C*. Instead, I have added a closure clause which specifies the absence of certain other factors. The further step I have not taken is unnecessary, in my view, for two reasons. First, the absence of destabilizing considerations (that is, the absence of reasons that weigh against the reasons one has for the commitment) will be sufficient to rule out implausible cases as legitimate instances of value commitments.

Second, the fact that I do not include "taking oneself to have reasons for one's reasons for having a commitment" in the analysis of the original commitment does not mean that people *will not* have this additional commitment; it simply means that we can understand the original value commitment without attributing these additional attitudes to agents. In fact, people will tend to have this additional commitment because there is pressure on us to justify the reasons for our commitments when we engage in deep deliberation. As we shall see in the next section, as deliberators we are concerned to have a certain kind of coherence in our values, and this means that we will tend also to be committed to the values indicated by the reasons we have for our commitments.

Now that the stable attitude account has been set out, it will be useful to point out the ways in which it differs from the view that our desires have normative force in and of themselves. On the stable attitude account, not everything a person is motivated to pursue will count as normative.

First, notice that on this account a person who takes her attitude to be supported by *some* reasons will not necessarily count as having a value commitment. Someone could have a pro-attitude toward having reflection on considerations C sustain her attitude A, and she could have an intention to continue to be so disposed, and yet not count as having a value commitment if it were the case that whenever she thought about a different set of considerations, she no longer cared about it. For example, imagine someone who loves gambling and has an intention to pursue it which is stable upon her reflection on the adrenaline rush she gets when she gambles. She has the relevant pro-attitude toward her commitment being stable in the light of considerations of this kind; that is, she thinks the fact that an activity gives her this rush is a good reason for having a commitment toward it. However, whenever she thinks about the time and money that gambling takes away from her other pursuits, or the fact that her desire to gamble is the result of an addiction, her intention to continue gambling wavers. The gambler departs from the paradigm case of a value commitment insofar as her intention to maintain her commitment to gambling is destabilized by reflection on other considerations about it. In paradigm cases of value commitments, destabilizing reasons against one's commitments are absent.

Secondly, these destabilizing reasons must be ones the agent regards as *good* reasons not to have the value commitment she has. It cannot be that a person does not count as having a value commitment if there is just *any* further set of considerations reflection on which would disrupt her attitudes. The gambler might lose her motivation to gamble whenever she reflects on death because such reflection depresses all her motivations, but this does not seem to diminish the normative significance of her attitude by her own lights. What must not be the case if a person is to count as having a value commitment is that there are some destabilizing considerations that are such that the person has a pro-attitude toward states not remaining stable in the light of these considerations. In other words, if the gambler's commitment to gambling is to count as an ideal case of a value commitment, it must be the case that there are no destabilizing considerations that she regards as good reasons not to be committed to it.

To summarize, let us return to an earlier example. Antonia's commitment to independence, is constituted by a pro-attitude toward independence and an intention to continue to have this attitude, where this pro-attitude remains stable upon reflection on the considerations she takes to be the reasons for remaining committed to independence. She approves of states of herself remaining stable upon this kind of reflection, and there are no considerations such that she approves of her commitments becoming unstable upon reflection on considerations of this kind, and intends to continue to approve of this pattern of destabilization. Notice

that the stable attitude account provides a compelling explanation for why some desires do not give us normative reasons for action. The fleeting desire to jump off of a tall building, for example, does not provide us with reasons because it is very much detached from other attitudes and commitments we have. In other words, this desire has no normative authority because it is not part of an appropriate pattern of attitudes.

In the preceding sections I have argued, by appeal to paradigm cases of value commitments, that to have such a commitment is to be in these mental states. When one has this particular arrangement of mental states, one counts as *valuing* something rather than merely wanting it. Our value commitments are normative for us in the sense that they provide us with reasons that justify our actions. The strength of these reasons depends on the stability and coherence of the pattern of attitudes that constitute our value commitment. Normativity, then, consists in a relationship between a person's attitudes and her ends: a person is related to her end in such a way that she has normative or justifying reasons when she has the set of attitudes toward that end specified by the stable attitude account. A person can take herself to have normative reasons when she does not actually have any if she is mistaken about having the appropriate pattern of attitudes.

There is one major obstacle to describing normativity in this way, however. So far, we have not seen how value commitments and our reasons for having them are subject to justification, improvement, or correction. Value commitments, as opposed to simple desires, are supposed to have the authority to justify what we do, but they could not have this authority unless they had some claim to correctness. If our value commitments are to have the authority they need, it cannot be the case that whatever psychological states we happen to have are as good as any other. There must be standards for improving the psychological states with which we begin.

4. Standards for value commitments: why it is not the case that anything goes

Our commitments can be relatively stable, and relatively well-informed without being ultimately stable and informed. If my account of the normativity of these commitments is accepted, the commitments with the most normative authority are those that most closely resemble the paradigm case in which the pattern of attitudes has ultimate, rather than merely temporary, stability. The paradigm case of a value commitment, therefore, serves as a standard for evaluating the relatively stable commitments we have now. The claim that the standards that arise from the stable attitude account of value commitment are standards for *improvement*

will be confirmed if the changes warranted by these standards speak to the deep concerns we have when we deliberate about our values. I elaborate these points in the discussion that follows.

To see how the features of our value commitments provide standards of evaluation for our ends, we can focus on cases in which changes of value commitments are warranted on the stable attitude account. On my view, it can be rational for an individual to change her value commitments because she can improve the stability of her set of attitudes in three ways. First of all, she might be ignorant of some destabilizing considerations that are such that she would have a favorable attitude toward reflection on these considerations destabilizing her commitment if she were aware of them, but as yet she has no attitude toward considerations of this kind. Through deliberation she might discover reasons not to have a commitment she previously thought was justified. It is often through deliberation that we find out about new considerations of this kind, that is, considerations that destabilize our value commitments in ways we endorse. In this case the person's value commitment is unstable insofar as it depends on a crucial lack of information.²² She can improve her value commitments by eliminating this source of instability through deliberation.

Secondly, a person might find in deliberation that her attitudes, although part of a relatively stable pattern of attitudes, do not withstand the closest scrutiny. She might find that the reasons she had for her commitment were not ones she can ultimately accept as justifying reasons, when she has more information or has reflected more thoroughly. In other words, upon examining her reasons she might find that she does not have a stable pro-attitude toward commitments remaining stable upon reflection on these reasons. This might be because she discovers considerations that destabilize this pro-attitude, or it might be that she can find nothing about her reasons that is worthy of approval. Finally, she might find that there is a new option, a possible end she had not previously considered, that is better supported by the good reasons she does have.

Silas Marner makes a profound change in his value commitments, and we can see how this change is warranted on the stable attitude account. First of all, Silas seems not to have thought of the way that his commitment to his coins encouraged his loneliness. He was therefore unaware of considerations that might have destabilized his attitude had he reflected on them. Secondly, Silas learns, through his experience with Eppie, that there are other commitments that can bring him more happiness than his commitment to his coins. In this way he comes to see that the reasons for loving his coins form a more stable pattern of attitudes in support of a different commitment. Eliot describes Silas's conversion in the following passage:

The gold had kept his thoughts in an ever-repeated circle, leading to nothing beyond itself; but Eppie was an object compacted of changes and hopes that forced his thoughts onward, and carried them far away from their old eager pacing towards the same blank limit – carried them away to the new things that would come with the coming years, when Eppie would have learned to understand how her father Silas cared for her; and made him look for images of that time in the ties and charities that bound together the families of his neighbours. The gold had asked that he should sit weaving longer and longer, deafened and blinded more and more to all things except the monotony of his loom and the repetition of his web; but Eppie called him away from his weaving, and made him think all its pauses a holiday, reawakening his sense with her fresh life, even to the old winter-flies that came crawling forth in the early spring sunshine, and warming him into joy because *she* had joy.²³

Until this point, Silas has not reflected on the way in which his love of his coins has caused him to work all the time and forgo other pleasures he might have enjoyed. That Eppie brings him the kind of joy that the coins never did, and that his commitment to Eppie brings him new experiences and closer relations to other people, are considerations that Silas now reflects on in such a way that the more stable pattern of attitudes is one in which he devotes himself to Eppie rather than to his coins.

A person's reflection on her commitments can be better or worse according to the above standards. She may ignore many destabilizing considerations because she is deceiving herself about their significance. She may overlook the fact that the reasons she has do not really support her commitment, or she may overestimate the instability of one of her commitments because of a fear that things will not work out as she hopes. Finally, she may neglect to consider other options that would leave her with ultimately more stable commitments. Moreover, we can see that changing our values on the basis of new information or insight about one's circumstances, options, and psychological states is the kind of change that answers to deep concerns we have in deliberation. These kinds of changes meet our concern to discover which goals are worth pursuing, or which ends worth valuing, given those features of ourselves that we find important on reflection, and to choose our goals and values in such a way that our choices can bear our reflective scrutiny.

We can now see the role of choice in the stable attitude account. Deliberators cannot simply observe which of their attitudes are most stable and rely on these to determine what is normative for them. Because our attitudes can be recalcitrant, they do not always correspond to what we discover in deliberation would be the most stable pattern. Therefore, we must sometimes choose to try to develop new attitudes in light of what we learn in reflection. Such choices will never be arbitrary, however, because they are made for the purpose of conforming to what we discover in deliberation about our attitudes, the world, and the relationship between them.

On the view I have defended, the existence of normative force is explained by patterns of attitudes, rather than by the existence of rational principles. The fact that normativity (or, our having justifying reasons for our actions) is explained by patterns of attitudes does not mean that all psychological states are equal and that therefore anything goes. There are standards for evaluating our commitments, which allow us to make judgments about what we have best reason to do. Despite this advantage of the stable attitude account over the view that every desire is normative by itself, my account will give rise to some objections. I turn to these objections in the next section.

5. *Problems for the stable attitude account*

I have argued that we take our ends to be normative when our attitudes stand in a particular relationship to these ends. We experience justifying reasons, then, when we have a genuine value commitment and we have standards of evaluation for the attitudes that make up this commitment. In this section I take up two of the main reasons why one might doubt that I have characterized our experience of normativity in the right way.

5.1 HOW CAN COMPLICATING THE HUMEAN PICTURE HELP?

One might say that the standards I have articulated for evaluating which commitments count as normative do not go beyond our own attitudes and that therefore this account is no more satisfying than the account which says that our desires are normative all by themselves. If there is nothing normatively significant about my having one pro-attitude, then adding more attitudes to the mix is not going to help.²⁴ Thus the thought here is that the stable attitude account does not adequately capture our experience of normativity because we simply cannot take motivational states to be reasons, no matter how they are arranged.

Whether complicating the picture helps explain how our value commitments are normative depends on the reason for which there is nothing normatively significant about my simply having some solitary pro-attitude. I suggest that the reason one pro-attitude by itself has no normative force has to do with the fact that we can take these pro-attitudes to be misguided, ill-informed, unfortunate, or sick. In Kantian terms, we have the capacity to stand apart from our desires and to evaluate which ones it makes sense to act on, which ones are worth endorsing, and which ones are not.

But it is not so obvious that we can judge our value commitments, as I have described them, to be misguided, ill-informed, unfortunate, or sick because if we did so, then there would be destabilizing reasons that we

would regard as good reasons not to have those commitments. In other words, value commitments, on the stable attitude account, are such that we could have no reason to judge that they are bad because it is an integral part of the account that an ultimately stable commitment (a commitment that we can take to be normative without need of revision) is one for which there is nothing we regard as a good reason not to have it.

In light of this fact, to maintain that we cannot take patterns of attitudes to have normative force seems to be mere bias against the Humean position. Now this kind of bias might be justified if it were the case that in order to have some standards of evaluation for our commitments, we must posit attitudes that are beyond critical evaluation. A position which posits unrevisable attitudes is not acceptable because it would have no explanation for the normative force of the uncriticizable attitudes. But the stable attitude account does not privilege any particular attitudes in this way. Although we could not reflect critically on all of our attitudes at once (because some must remain stable to provide standards of evaluation for the ones under scrutiny), and although some of our attitudes have such a central role in our patterns of value commitments that we rarely feel the need to reflect on them, the stable attitude account holds that all of our attitudes are in principle within the scope of critical reflection.

5.2 UNIVERSALITY

The stable attitude account provides standards for evaluating our commitments, but these standards are not entirely independent of the other commitments a particular deliberator has. This means that the stable attitude account does not guarantee that different deliberators will reach the same conclusions about what has value, and many will find this unsettling. If we are not assured that there will be the same standards for everyone, the objection goes, then we do not have real standards at all; our experience of normativity tells us that such things as values and justifying reasons must be universally binding.

This kind of objection is most pressing when we think of the case of commitments to moral values. Notice, though, that if we think of other kinds of value commitments, this problem is much less troubling. Some of the commitments we have are to values that are, we might say, personal rather than moral. Notice that people do tend to think that such commitments are normative, and that we do not take the fact that there are no standards which require everyone to value our personal projects to count against their normativity. We do think that our particular projects and relationships (such as a commitment to a career, or a hobby such as playing music) have value, that we ought to spend time and energy in pursuit of these commitments, and that these commitments provide us with reasons to do things. However, we do not think that everyone ought

to be committed to, say, being a philosopher or playing the violin. The standards we use to evaluate these commitments as appropriate for us do not make it the case that these projects must be valued by everyone. Reflecting on the non-moral case shows us that it makes sense in some cases to think that there are standards for evaluating the normative force of our commitments that are not universally binding. Our experience of normativity is not uniformly an experience of universal values.

I cannot fully defend the extension of the stable attitude account to the moral case here, but I do want to make a few suggestions about how this concern can be deflected. To do so it will help to have in mind some particular moral values that are frequently thought of as reasons for having various moral commitments. The fact that some practice causes suffering is taken by many people as a reason in favor of a commitment to protesting or avoiding this practice. For example, many ethical vegetarians take the suffering caused to animals as their primary reason not to eat meat.

Now, on the stable attitude account, taking suffering as a reason amounts to having a favorable attitude toward one's commitments remaining stable upon reflection on the fact that this commitment prevents (or at least does not participate in) suffering. To have normative force, these commitments must remain in their stable patterns upon reflective scrutiny. Here we can see clearly the source of concern about a Humean account: the fact that the commitment to vegetarianism, for example, has normative force seems to depend on the agent having a very specific, stable pattern of attitudes. Other agents need not have this pattern and this means that for them vegetarianism may have nothing to recommend it.

But when we focus on the kinds of moral reasons on which we typically rely we can also see that things are not as dire as they may have seemed at first. This is because the basic sorts of moral reasons we offer to justify our moral commitments *are* widely shared; they are typically commitments that are deeply implicated in people's patterns of value commitments, and ones we cannot very easily change. Further, not only are commitments to moral values deep and widely shared, but there are significant social pressures that work to maintain these values.²⁵

Even with the above observations in mind, one might not find the stable attitude account of the normativity of *moral* commitments satisfying. One might insist that we cannot live with the fact that the normative force of our commitments ultimately depends on our having certain attitudes, and that if this is what normativity turns out to be, then it is does not live up to our experience of it. Our experience of moral values and reasons is that they are non-optional; their normative force does not depend on our having any particular set of pro-attitudes. I will not dispute that we experience moral values in this way. The question, though, is whether a Humean account such as this one can make sense of this experience. I think there is some reason to think that it can.

Notice first of all that living with some form of Humeanism is not the same as living with existential heroism. It is not merely my free choice that makes my commitments valuable. Rather, my commitments are normative if I can reflect on them and see that they are stable under scrutiny, and, in fact, it is not the case that I can choose to change just any of the attitudes I have. This scrutiny involves consideration of my reasons for having the commitments I have, and of the relationships between my various commitments, reasons, and attitudes. Reflective scrutiny, on my view, does not result in discovering a law of reason, but it does result in examining the reasons one has in light of the facts, and in coming to develop commitments that are supported by the most refined set of reasons one can find.

Secondly, the stable attitude account does not entail any particular view about how we ought to regard the commitments of others. Our attitudes toward other people's values will, in some cases, constitute additional members of our set of value commitments, and these commitments too will be analyzed in terms of the stable attitude account. It might be that when it comes to moral values, we always disapprove of others who do not share our values because we think there are reasons for our values that these others fail to appreciate. If this is so, we can understand this disapproval as a commitment to the universality of the moral value in question. Our approval of those who agree with us, and disapproval of those who do not, are the relevant constitutive attitudes of this commitment. Our taking there to be reasons for this approval (or disapproval) can in turn be analyzed in terms of pro-attitudes toward attitudes remaining stable upon reflection on certain considerations. The experience of normativity is an experience we have when our attitudes are related to our ends in a certain way; saying this does not commit us to the view that our attitudes toward other people's values must be permissive.²⁶

6. *Conclusion*

I have argued that value commitments can be understood in terms of patterns of attitudes. In particular, to have a value commitment is to have a pro-attitude, an intention to continue having that attitude which constrains future deliberation, and a further pro-attitude toward your attitude remaining stable upon reflection on certain considerations about the commitment, and for it to be the case that this pattern of attitudes is not destabilized in ways approved of by the valuer. In virtue of these commitments, and the fact that there are standards for improving them, we do have normative or justifying reasons to pursue our ends. These reasons can be better or worse, depending on how well the value commitment from which they derive meets the standards described by the paradigm case.

There are some disadvantages of a Humean account. In particular, some will think that the Humean account does not give us precisely what we sought, especially when it comes to moral values. Whether we can live with this kind of account will depend, in part, on what the alternatives are. If the Kantian project is ultimately unsuccessful, and normative nihilism is our alternative, we might be more willing to continue as usual in our normative practice while accepting that this practice is not underwritten by pure principles of practical reason. For some of us, avoiding nihilism may seem to require a certain amount of heroism, but such people may very well agree that a little "Humean heroism" is worth the price of escaping the desperate position that there are no values and no reasons to do anything.²⁷

Department of Philosophy
University of Minnesota

NOTES

¹ This position is known as internalism. See Williams, "Internal and External Reasons" in his *Moral Luck* (Cambridge: Cambridge University Press, 1981).

² For relevant discussions see Simon Blackburn, "Practical Tortoise Raising" *Mind*, 104 (1995), 695–711; James Dreier, "Humean Doubts about the Practical Justification of Morality," in Garrett Cullity and Berys Gaut (eds), *Ethics and Practical Reason* (Oxford: Clarendon Press; New York: Oxford University Press, 1997); Jean Hampton, "On Instrumental Reason," in J. B. Schneewind (ed.), *Reason, Ethics and Society: Themes from Kurt Baier, with his Responses* (Chicago: Open Court, 1996); Christine Korsgaard, "The Normativity of Instrumental Reason," in Cullity and Gaut (*op. cit.*); Peter Railton, "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action," in Cullity and Gaut (*op. cit.*).

³ *Op. cit.* Korsgaard makes a similar argument in *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996).

⁴ "The Normativity of Instrumental Reason," *op. cit.*, 253.

⁵ This is undoubtedly a caricature of existentialism. My purpose in discussing it is to motivate the need to develop a view that is like heroic existentialism in certain respects.

⁶ For a related discussion of the idea that first order desires cannot represent our will, see Frankfurt, "Freedom of the Will and the Concept of a Person" and "Identification and Wholeheartedness," in his *The Importance of What We Care About* (Cambridge: Cambridge University Press, 1988). For an account that defines what it is to value in terms of a simple desire see Donald Davidson, "Intending" in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), 86. For criticisms of this view see Gary Watson, "Free Agency," in his *Free Will* (Oxford: Oxford University Press, 1982), and Michael Smith, *The Moral Problem* (Oxford: Blackwell Publishers, 1994).

⁷ By "pro-attitude" I mean any positive affective state. Such states can include desires, emotions, and other feelings, and I take them to be dispositional, not occurrent.

⁸ In my discussion and defense of the stable attitude account of value commitments I emphasize our experience of normativity because this seems a good place to start if we are trying to understand such normative notions as justifying reasons and values. Of course,

not everyone would agree that this is the place to start, and some would even argue that our experience of the normative dimension has no bearing on what actually has value, or what our reasons actually are. The present paper is not an argument against this view, but, I hope, a demonstration that we are at least not forced to take this kind of radical stance.

⁹ For further discussion and support of the point that our desire do not give us normative reasons by themselves see Frankfurt, "Freedom and the Concept of a Person," *op. cit.*; and Scanlon, *What We Owe to Others* (Cambridge, Mass.: Belknap Press of Harvard University Press, 1998). One might think that desires do give us reasons, but that these reasons are defeasible. However, it is consistent with common sense to think that some desires, such as Watson's case of the frustrated mother's desire to throw her baby out the window, provide us with no reason to act at all (Watson, "Free Agency," *op. cit.*, 100–101). Therefore, it seems that the position that all desires provide defeasible reasons should be resorted to only in case no other explanation of normativity can be defended. I thank David Schmidtz for pressing me on this point.

¹⁰ George Eliot, *Silas Marner* (New York: Bantam Books, 1992), 18–19.

¹¹ *Ibid.*, 116.

¹² This is not to say that Antonia *cannot* deliberate about the value of independence. In fact, it might be quite natural for her to do so if independence comes into conflict with another central value for her. The point here is that in the ordinary course of events, when she is reasoning about what actions to take in the relatively short term, the value of her central commitments remains stable.

¹³ What I have said here about intentions as a feature of value commitments is informed by Michael Bratman's account of intention in his *Intention, Plans and Practical Reason* (Cambridge, Mass.: Harvard University Press, 1987). I discuss his account in more detail in the next section.

¹⁴ Nor is it the case that the reasons we have for our commitments must be of a particular sort (such as instrumental or self-interested reasons). In some cases we might have a commitment to something because we think that it is valuable in virtue of its intrinsic features or nature.

¹⁵ This is not to say that we must always take our efforts to be justified by some further end. It might be that we take the pursuit of the end itself to be worthwhile. But even in such cases the justification we give ourselves goes beyond "I want it."

¹⁶ "The Normativity of Instrumental Reason," *op. cit.*, 250.

¹⁷ It seems that we can imagine a person who has great but fleeting conviction in the value of her commitments, for example, a person who floats from one project to the next without losing faith in the idea that these fleeting commitments provide real guidance. There are two things to say about such counterexamples. First, people who seem to follow this pattern often take themselves to be committed to the value of spontaneity or "living for the moment." If this is true, then these people have different, but nevertheless stable, value commitments. Secondly, if there are people who genuinely have no stable commitments (not even to spontaneity), it seems to me that this pattern is very hard to sustain over a lifetime. Such a person may very well become plagued with skeptical doubts about the value of any of their commitments. Notice that one can take oneself to have reasons even for ephemeral commitments and that therefore, stability is a distinct feature of our value commitments rather than a product of the judgment that there are reasons to have the commitment.

¹⁸ Here I draw on Michael Bratman's account of intention in his *Intention, Plans and Practical Reason*, *op. cit.*, and his "What is Intention?," Center for the Study of Language and Information, Report No. CSLI-87-69, (January, 1987).

¹⁹ In this way the account of intention I am using departs from one of the uses of "intention" in ordinary language.

²⁰ In emphasizing the importance of this feature, the stable attitude account is similar to Gary Watson's account of a person's values, although I give a different analysis of what it is to take oneself to be justified. According to Watson, "an agent's values consist in those principles and ends which he – in a cool and non-self-deceptive moment – articulates as definitive of the good, fulfilling, and defensible life" ("Free Agency", *op. cit.*, 105).

²¹ Recall that an intention is a state that functions to constrain practical reasoning by limiting the options amongst which one chooses and by constraining what could count as an acceptable option. In this case, I am talking about an intention to continue to have reflection on reasons sustain one's value commitments. Such an intention might constrain future practical reasoning by discouraging us from reconsidering all of our reasons for a particular commitment when we are deliberating about how best to fulfill that commitment.

²² The problem here might not be that the person is missing the facts. She may be missing an appropriate appreciation of the facts, or, she may have failed to draw the proper inferences from these facts.

²³ *Silas Marner, op. cit.*, 127–128.

²⁴ For treatment of a similar problem see Frankfurt, "Identification and Wholeheartedness," *op. cit.*

²⁵ See John Stuart Mill's discussion of the sanction of morality in chapter III of *Utilitarianism*, ed. John Gray (Oxford: Oxford University Press, 1991). See also, Dale Jamieson, "Animal Liberation is an Environmental Ethic," *Environmental Values* 7 (1998), 41–57.

²⁶ Simon Blackburn and Allan Gibbard treat the objectivity of moral judgments in a similar way. See Blackburn *Spreading the Word* (Oxford: Oxford University Press, 1984); and Gibbard *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, Mass.: Harvard University Press, 1990).

²⁷ For helpful comments on earlier drafts of this paper I would like to thank Simon Blackburn, Thomas E. Hill Jr., Geoffrey Sayre-McCord, David Schmitz, Michael Smith, John D. Walker, and the 4M discussion group. I would also like to thank the Office of the Vice President for Research of the University of Minnesota for support of my research on this paper.