

CHAPTER FOUR
THE BOTTOM-END THEORY OF NORMATIVE REASONS

4.1 Introduction

One aim of this chapter will be to argue against the view that normative reasons are grounded in the desires of the agent. Given (a) the fact that this view figures prominently in Korsgaard's and Frankfurt's accounts of the relationship we have been exploring, and (b) the large impact that this thesis' truth or falsity will have on any attempt to work out a connection between practical reasons and self-conceptions, I have chosen to give this issue special attention. Accordingly, this chapter will consist in an argument for the view that normative reasons for action are value-based (VBR) rather than desire-based (DBR), and to a lesser extent, for the view that reasons are fact-based (FBR) rather than belief-based (BBR). In the second half of Chapter 5, I will first explore how a commitment to VBR and FBR counts against Korsgaard's and Frankfurt's views, and then develop a VBR- and FBR-informed view of the reasons/self-conceptions relationship.

In terms of strategy, my approach to defending VBR will differ in important respects from the approaches employed by those who have come before me. First, unlike many traditional attacks on DBR which try to show that all desires are responses to values,¹ my focus will be on the first-person deliberative pressures that we face when thinking about what to do. Our best defense of VBR, I will argue, lies not in scrutinizing the mechanics of psychological states, but rather in our reflective need to justify our actions to ourselves and others. Second, while some VBR

¹ Many of these more traditional arguments owe their inheritance to the earlier work of Thomas Nagel's *The Possibility of Altruism*, Chapter 5 (for a more recent discussion, see Joseph Raz's "Incommensurability and Agency"). Somewhat ironically, the approach that I favor has been indirectly influenced by some of Nagel's more recent work on the pressure that we find ourselves under to adopt a more objective point of view in our thinking.

proponents have offered the beginnings of an argument of this sort (by discussing anecdotal cases in which we appear to take our reasons to be tied to values rather than desires),² I will try to parlay these anecdotal cases into a more systematic treatment of this topic that centers around a platitude about the way that reasons justify actions.

The argument that I will offer consists of two parts. First, after laying out some terminology (Section 4.2), I will raise the following question. Reflecting in a cool hour, what sorts of considerations do we—operating from the first person deliberative perspective—ultimately accept as counting in favor of performing certain actions (Section 4.3)? The answer I will suggest is: desire-independent values. More specifically, I will argue that the only justifications that we ultimately find satisfying are ones that bottom out in appeals to desire-independent values (hence the “bottom end” theory of normative reasons). Since our very concept of a normative reason for action is that of a pro-consideration that justifies action, this is an important conclusion. For if it is values that justify a given course of action, then it will be VBR—and not DBR—that better cashes out this platitude about the justificatory role of normative reasons.

The second part of the argument brings us to a second question—one that has been sorely neglected in the literature, even by VBR proponents who are sympathetic to the sorts of argument that I will offer. This second question is: does the mere fact that we take our reasons to be tied to values *itself* vindicate VBR (Section 4.4)? Since there are a number of philosophers who accept the phenomenological considerations that I’ll discuss, but who nonetheless reject VBR, this is a question we must address. In response, I’ll show that a failure to take these phenomenological considerations seriously would constitute an unnecessary act of violence against our very concept of a

² For example, despite the fact that T. M. Scanlon in *What We Owe To Each Other* is ostensibly arguing for a version of reasons externalism rather than VBR, he in effect offers an argument for VBR (or at least a close relative of it) by relying on anecdotal cases of this sort. One of the central strands of Jonathan Dancy’s argument against DBR in *Practical Reality* relies on a similar style of argument.

normative reason for action. As at least some of the values that we take ourselves to be tracking do in fact exist, it would be inappropriate to offer an account of normative reasons that isn't true to this phenomenological commitment of ours. I conclude by arguing for the view that normative reasons are fact-based rather than belief-based, which serves an apt (and necessary) compliment to VBR (Section 4.5).

4.2 Terminology

I have already discussed the difference between motivating and normative reasons, both in Chapter 1, and in passing in the previous two chapters. Given present purposes, however, it would be useful to reiterate the important fact that normative reasons *justify* certain courses of action (where the justification in question is of a non-excusing sort).³ This idea is so fundamental to our concept of a normative reason that we could posit it as a platitude. To claim that a normative reason for action exists is to claim that a certain action would be justified. The importance of this platitude for the arguments that follow will become clear shortly.

With this in mind, consider the following two claims, both of which purport to provide existence conditions for normative reasons for action:

- Desire-Based Reasons Claim (DBR): When a normative reason for me to ϕ exists, it does so (a) because ϕ -ing will satisfy a desire of mine (or the desire of a more rational version of myself), and (b) not because ϕ -ing would promote or realize desire-independent values (where a “desire-independent value” is a value the content and existence conditions of which obtain independently of the desires of any agent). This claim accordingly includes two theses about the existence conditions for an agent to have normative reason to pursue a given object: the

³ Recall that motivating reasons are sometimes offered to “justify” actions by excusing them (“Don’t blame me. The bottle said ‘gin’ not ‘petrol;’ I thought I was preparing the run-of-the-mill gin and tonic that he asked for”). Normative reasons don’t justify actions by excusing them in this way.

agent's caring about the object accounts for the existence of the reason, whereas the desire-independent value of the object (if there is such a thing) does not.⁴

- Value-Based Reasons Claim (VBR): When a normative reason for me to ϕ exists, it does so at least in part because a given desire-independent value exists. In other words, it is a necessary condition on R's being a normative reason for me to ϕ that ϕ -ing possess value-making or value-constituting properties that obtain independent of whether I or a more rational version of myself care about them.

The main issue here concerns the relationship between practical reasons and desires. DBR posits a very tight connection between practical reasons and the desires of an agent (such that the latter ground the former), and denies any connection between normative reasons and desire-independent values (such that the latter—if they exist—have no bearing on the former). According to VBR, on the other hand, there is no such connection between normative reasons and the desires of the agent (actual or hypothetical). VBR doesn't hold that facts about reasons and desires are unrelated, but rather that the existence of the former isn't determined by the existence of the latter.⁵ To illustrate, according to DBR, whether I have a pro tanto normative reason to donate money to disaster relief will be determined by whether I have desires that donating money would satisfy. According to VBR, if I have a pro tanto normative reason to donate money, this will be because donating money would realize desire-

⁴ There are accordingly two positions that are consistent with DBR: (1) the view that there are no desire-independent values (i.e. that a given object is of value only because someone has a certain attitude toward it), and (2) the view that there are desire-independent values, but they don't give rise to normative reasons.

⁵ To mark out the landscape within which I am working, it might be helpful to note where some prominent philosophers fall with respect to this distinction. Proponents of VBR (or some close relative of it) include Jonathan Dancy, Joseph Raz, E. J. Bond, Derek Parfit, T.M. Scanlon, and Philippa Foot (at least in her more recent work, including *Natural Goodness*). Proponents of DBR (or some close relative of it) include Michael Smith, Bernard Williams, Harry Frankfurt, Simon Blackburn, and Christine Korsgaard (at least on one interpretation of her position in *The Sources of Normativity*).

independent values that exist regardless of whether I have any desires that a donation might “speak to.”

It’s worth mentioning that the DBR/VBR debate differs from the reasons internalism/externalism debate.⁶ A reasons internalist maintains that it is a necessary condition on my having a normative reason to ϕ that either I or a more rational version of myself could be motivated to ϕ ; a reasons externalist denies this. To claim that the presence of a desire to ϕ is a necessary condition for the existence of a reason is not to necessarily to claim that a reason for action must be *grounded* in a desire. It is consistent with reasons internalism, in other words, to claim both (1) that I might have a reason to ϕ in part because I could be motivated to ϕ , and yet (2) that ϕ -ing has something to be said for it apart from the fact that it would satisfy an antecedent desire of mine.⁷

4.3 The Phenomenology of Deliberation

My aim in this Section will be to show that, when we reflectively consider matters from the deliberative point of the view, it becomes clear that our concept of a normative reason for action is that of a pro-consideration grounded in desire-independent values. The argument for this conclusion involves the following two steps. (1) It is a conceptual platitude that normative reasons for action justify certain courses of action. (2) Upon reflection, it appears to us—deliberating from the first

⁶ The literature on this debate has grown quite large. For some good discussion of this issue, see Bernard Williams “Internal and External Reasons, with Postscript,” Robert Audi’s “Moral Judgment and Reasons for Action,” David Sobel’s “Explanation, Internalism, and Reasons for Action,” John Robertson’s “Internalism, Practical Reason, and Motivation,” Brad Hooker’s “Williams’ Argument Against External Reasons,” John McDowell’s “Might There Be External Reasons?,” and Scanlon’s appendix on Williams (Scanlon: 363-73).

⁷ VBR and reasons externalism are nonetheless close bedfellows. Considerations that prompt one to hold VBR will likely prompt one to hold reasons externalism as well (since the general intuitions that the normativity at issue is provided by values speaks equally against the view that desire-satisfaction or the bare possibility of being motivated to do something have much to do with what ultimately counts in favor of performing certain actions).

person point of view—that the only sorts of considerations that can ultimately justify actions are considerations involving desire-independent values (i.e. properties that strike us as “worth having” regardless of whether a given agent cares about them). I will proceed on the assumption that claim (1) is uncontroversial. In what follows, my focus will be accordingly on claim (2).

In our everyday deliberations, our decisions to undertake certain pursuits typically appear to us to be responses to the value of these pursuits. I don’t, for example, take myself to have reason to study philosophy simply because I want to; I take the study of philosophy to be a valuable pursuit, one that everyone has at least some reason to pursue (i.e. on account of the value of philosophical study, there is always at least a *pro tanto* reason counting in its favor). If it turned out that philosophy was not valuable in the way that I think it is—for example, if it turned out that this study took me further away from what is good and true in life, or if it didn’t actualize important capacities of mine—I might no longer take myself to have reason to be a philosopher. The reasons that I think I have for studying philosophy, in other words, don’t appear to me to be ultimately grounded in the fact that this pursuit answers to some desire that I happen to have.

In sketching this first example, I haven’t yet done anything entirely new; this is precisely the sort of illustration that other proponents of VBR have emphasized. It is, however, important that we say a bit more about the character of these sorts of phenomenological considerations. For surely many people will object that when they think about what to do, they don’t spend much time reflecting on desire-independent values. Many philosophers will insist that a survey of desire-independent values wasn’t what brought them to philosophy in the first place.

In response to concerns of this sort, it is crucial that we think of VBR less in terms of the statistical frequency with which we *actually* worry about values in our

deliberations, but rather think of it in terms of the sorts of considered justifications that we are *ultimately pressed to offer* for our practical pursuits. Put in these terms, we can restate the underlying phenomenological point as follows: to justify fully my commitment to a given project to either myself or to others, it seems (once again, from my point of view as a deliberating agent) as though I must ultimately connect the project in question to desire-independent values. “Because I simply felt myself drawn to this project” will rarely suffice as a justification for many of our pursuits, either in our own eyes or in the eyes of others. Even if you didn’t explicitly take up philosophy in response to perceived desire-independent values, when faced with either a pushy interlocutor or nagging doubts of your own concerning your career choice, the sort of justification you will be driven to offer will be one that rests on claims about the desire-independent value of philosophy. I accordingly assume that most of us will be quite unsettled by the thought that the deepest justification we can give for our life’s pursuit is: “I’m a philosopher because I’m interested in philosophy, and this interest isn’t a response to values of any sort (to include the value of enjoyment). I simply find myself drawn to philosophy, and there is no further justification to be given.”⁸

This is not to say either (1) that we don’t in fact have some desire to pursue the projects that engage us, or even (2) that our desires can play no role in the final determination of what we think we ought to do. The claim is rather that—from our own points of view as deliberating agents—the fact that a given project would satisfy a desire is not a necessary condition on our having at least a *pro tanto* reason to pursue that project, nor can such a desire ever be a sufficient condition for such reasons.

To further illustrate this point, consider the following exchange, which is based on an example provided by Warren Quinn:

⁸ For a more detailed discussion of the practical and psychological problems that arise if one believes that there are no objective values for one to track, see David Wiggins’ “Truth, Invention, and the Meaning of Life.”

Dick: “I have good reason to turn on radios whenever I pass them simply because I have a strong desire to do so.”

Jane: “I don’t understand. Is this because of some superstitious belief on your part? Do you think something bad will happen if you fail to turn on radios that you come across?”

Dick: “No. I just want to turn on radios whenever possible, and hence I have a perfectly intelligible reason to do so.”

Jane: “I still don’t understand. Is this because of the pleasure you get out of turning radios on, or the pain you would get if you refrained?”

Dick: “No, pleasure and pain have nothing to do with it. I just want to turn the radios on, and that’s all there is to it.”

I for one don’t see how Dick could either expect anyone to accept this sort of justification, nor be sincerely satisfied with it himself. As long as Dick and Jane are willing to reflect on the matter, I don’t see how either of them could be fully satisfied with a justification that doesn’t eventually bottom out in desire-independent values.⁹

It’s important to consider bizarre cases like this one because it is easy to think that desires do all the justification we need in some cases (like choosing to be a philosopher) simply because we are assuming that the objects in question are valuable (such that “I want to study philosophy” seems sufficient to justify the project simply because we *already* see philosophy as worthy of pursuit). In instances like the radio case, however, when we don’t assume that there is some value in the proposed action, we tend to think that the desires in question are less convincing as justificatory devices.

⁹ Berys Gaut discusses another example of this sort that could be used to marshal similar intuitions at (Gaut: 178).

We can also use cases of this sort to illustrate precisely what one is committed to if one accepts DBR. For any case in which an agent proposes doing something odd, the DBR proponent has to admit that the mere fact that the agent wants to perform the action constitutes all the justification one could need. Imagine an agent claiming that she really wants to cut off her legs. Assume that this desire isn't a response to any consideration of value, and that this agent has carefully thought out how the removal of her legs will impact her other plans. While the DBR proponent may think this is an odd desire, once it is acknowledged that the agent both has and welcomes this desire (even after deep reflection), there should no longer be anything mysterious about her proposal. By the DBR proponent's lights, this agent has offered as compelling a justification as one could hope for.

Bear in mind that my central claim here is about what *normative* reasons look like from the deliberative point of view, since doing so will help us avoid two worries. First, this central claim is not invalidated by the fact that people sometimes do end up choosing a given option "just because they want to." Sometimes agents aren't concerned with justification at all—they just "want to do what they want to do." Since the agents in such cases aren't really worried about normative reasons for acting in the first place, these cases aren't counterexamples to my argument.

Second, the mere fact that some people don't seem to face the reflective pressures that I've been discussing is no threat to this account either. The point is (1) we can always bring agents to face these reflective pressures by demanding that they justify their reason-claims (agents who have been trained to stubbornly resist such pressures are covered by my first point above), and (2) once we have got these agents to adopt a more reflective standpoint, they will tend to be troubled if they cannot provide a justification that bottoms out in appeals to value. Agents typically respond

as though this more reflective standpoint had a certain authority as the final arbiter of what is reason-giving.

This point about the reflective pressures that we are under might also help show that VBR is validated even in the case of pursuits that seem fundamentally desire-dependent. Take the case of hobbies. It is indeed true that few of us take up hobbies by surveying the world of desire-independent values, but faced with reflective worries about whether our hobbies are a good use of our time, the realization that there are no desire-independent values associated with our hobbies might lead us to abandon them. My interest in single malt Scotch may not have initially developed as a response to desire-independent values, but when I slip into a reflective mood and begin to wonder whether taking an interest in drinking pungent liquids is silly, values are what I look for to settle this issue. I might think, for example, that there is desire-independent value in (1) developing the discriminating powers of one's senses to the point where one can appreciate geographical influences on a Scotch's development, and (2) learning about the rich history of how cultural changes have affected the way Scotch is produced and marketed. When I become convinced that these considerations represent genuine tokens of desire-independent values, I begin to think that this hobby really does have something to be said for it.

It is important to note that whether these examples of purported values involve *actual* values is completely irrelevant. Even if I am mistaken about the value of developing the discriminating power of one's sense organs, this is no threat to my central claim. The point is simply that we face, or can be brought to face, reflective pressures that prompt us to pursue only projects that are associated with desire-independent values.

One might worry that it is at precisely this point that VBR shows itself to be unpalatable. For if this view is correct, it would seem that many people might have to

give up their hobbies. I've recently heard of a man who collects stamps from around the world featuring bagpipes. Can we really point to any general categories of value that this pursuit instantiates? As it turns out, I think there is at least one important category of value that might underwrite this particular hobby, and moreover, one that we can all appreciate. On the assumption that enjoyment is a desire-independent value, then one might cite this value as the basis of a *pro tanto* reason for this stamp enthusiast to go about his business.

Some may think that this is a rather ad hoc attempt to defuse what looks to be an embarrassing consequence of VBR, but I truly don't think it should be seen this way. The fact that someone enjoys something does appear to count in its favor. Perhaps even more importantly, we can give examples to show that enjoyment is a genuine desire-independent value; i.e. that it could justify a given course of action apart from whether the agent desires her own enjoyment. Think of cases in which you have tried to get a depressed friend outside to do something enjoyable even when your friend professes a disinterest in doing enjoyable things. When such attempts manage to draw someone out of a funk, it does indeed appear as though the depressed friend had reason to do something enjoyable even in the absence of any desire to do so.¹⁰

But perhaps the examples I've been providing so far aren't yet enough to convince everyone that all of our reasons are tied to desire-independent values. What about our choice in friends? Isn't it the interest we take in certain people that really grounds our reasons for treating them as friends?

Once again, I'm inclined to argue that bottom-end considerations of value can play a big role in justifying friendships. I don't take my friendship with Linda to be

¹⁰ One might reply that the appropriate desires in this sort of case would need to belong to a fully rational version of this agent, and so a DBR proponent can accommodate these intuitions. This might be the case, but I'm inclined to think that even if a fully rational version of this person (specifically of the sort that DBR proponents—like Michael Smith—discuss) didn't care for enjoyment, it isn't clear that this "rational" disinterest would prove that enjoyment isn't a legitimate value. To see why this might be so, see my discussion of the "Rationali case" below.

justified wholly by my desire to be around her. Linda possesses a number of qualities that I think are valuable. She has a good sense of humor, and is very kind and considerate toward others. If you asked me why I am Linda's friend, I would cite these qualities and expect you to recognize their value. If, however, it turned out that all of Linda's witty observations were cribbed from other people, and if I discovered that she enjoyed torturing small animals, I would no longer take myself to have much reason to treat her as a genuine friend (at least not in the way that I did before). Thus even in cases of friendship, one can argue that the phenomenology of deliberation supports VBR over DBR.¹¹ This example also further illustrates the usefulness of counterfactual testing as a general strategy for supporting VBR—even if in many cases the counterfactuals in question come off as a bit bizarre.¹²

To put this point more generally, our experience of error—of being wrong about what is of value and what is not—is further evidence in support of VBR, at least in certain contexts. Since we sometimes take ourselves to be mistaken about the worthiness of our projects, this shows that our interest in certain activities really does

¹¹ One might worry that this implies I should break things off with Linda as soon as someone more humorous and considerate comes along. In response, one could argue that long-standing relationships put one in touch with a number of values that would be lost if one threw off one's friends for "better models" (proceeding on the rather implausible assumption that I can't be friends with both Linda and this other person at the same time). My shared experiences with Linda, the fact that I can trust her and that she knows me in some ways better than I know myself are goods that would be lost if I ended my relationship with Linda. I don't mean to imply that one should view one's friends as mere containers of valuable characteristics, or that one should shop around for friends on the basis of the values they could potentially bring one. The idea is simply that we don't have the friendships we have for no good (desire-independent) reason.

¹² It has recently occurred to me that the repulsion many have towards VBR may be the result of a failure to distinguish between *de dicto* and *de re* interpretations of the sorts of claims I've been making. When I claim that we respond to values in our decision making, I don't mean to say that we make our decisions under this description. When I choose to keep Linda on as a friend, I do so in direct response to her value, not in response to the recognition that there are some values here, and since I am in the business of responding to values, it would make sense to keep her on as a friend. It isn't as if I see a sign hovering over Linda with the words "Value Here," and this is what I'm responding to. Understood in a *de re* sense, some of my claims may come off less mercenary.

track our assessment of their value.¹³ Consider rather straightforward cases in which the value in question is that of subjective satisfaction. Based on our own experiences, there are clear cases in which the object of our desires—and thus the satisfaction of our desires—turns out to disappoint our expectations. E.J. Bond has put this point well in discussing two senses in which we can speak of an agent’s “satisfaction” in achieving the object of her desire:

In one sense of ‘satisfy’, it is a sufficient condition of one’s *desire* being satisfied that one *gets* what one wants; if you get what you want, then (it follows) your want is satisfied. But it does not follow from this that *you* are satisfied in the sense that you derive any pleasure or satisfaction, in the other sense of *felt* satisfaction, from the acquisition, since it may yield no reward at all (Bond: 45).

Bond goes on to claim that even when we do find satisfaction in acquiring the object of our desires, “...the value [doesn’t] lie in the satisfaction of desire *per se*...the value...lies in the enjoyment of the thing, not in the fact that some desire has been satisfied” (Bond: 61). The upshot is that the mere presence of a desire—even a fairly reflective desire—isn’t sufficient to confer value upon an object even in cases in which subjective satisfaction is what counts in favor of the option, since we are quite familiar with cases in which it becomes clear after the fact that a desired object never really had anything counting in its favor in terms of satisfying *us*. It is rather the enjoyment that we get out of a given object that provides the pro tanto reasons.

Setting aside cases of this last sort, there are also cases in which we decide that certain pursuits weren’t really valuable, despite the fact that we were “satisfied” with them at the time. To give an example, there was a time in perhaps many of our lives when we thought that popularity (perhaps at the expense of the feelings of others) was

¹³ For discussion of the ways in which desires and the pleasure we take in activities might play an epistemic role in our efforts to discover what is of value, see Dennis Stampe’s “The Authority of Desire” and Elijah Millgram’s “Pleasure in Practical Reasoning.”

of great value. Looking back on these adolescent beliefs, we may be inclined to think of these beliefs as misguided. My present attitude toward this sort of popularity isn't explained simply by the fact that my desires have changed since my early teenage years—it's not just that I care about different things now—this attitude is explained rather by my belief that in my youth I had desired the *wrong things* (what I thought was of value was in fact not).

To this point I have been focusing on versions of DBR that link one's reasons for action to the actual desires of a given agent. In the interest of rounding out this defense of VBR, however, it might be useful to consider both a "rational desires" version of DBR, as well as a "hybrid" view that contains elements of both DBR and VBR.

Beginning with the former, the general line of argument that I've been offering holds with equal force against versions of DBR that tie one's reasons to the desires of fully rational agents. According to Michael Smith, for example, my normative reasons are determined by the advice that a fully rational version of myself would give me (where the advice in question is based on what this fully rational person desires that I do).¹⁴ To show where this version of DBR goes wrong, imagine that I am confronted with a fully rational version of myself (call him "Rationali"). While I might agree that Rationali has a good handle on the reasons that apply to me, it still seems as though these reasons will not be determined by his desires. This point becomes especially clear if Rationali gives me advice that I find counterintuitive, like "join the Taliban." Upon hearing this advice, I will surely want to know *why* Rationali desires that I do these things. And the idea again is that the only sort of explanation that Rationali can give that I'll find satisfying is one that bottoms out in desire-independent values.

¹⁴ See (Smith: Chapter 5).

A similar conclusion follows if we look at things from Rationali's point of view. If my arguments have been correct, Rationali won't take his own desires as such to ground any reasons for action; Rationali—just like all the less than fully rational people out there—will see reasons as grounded in desire-independent values. And if Rationali takes himself to be responding to desire-independent values, any reasons I have will be ultimately grounded in these values.

What about a hybrid view, according to which some of our reasons are determined by objective values and some are determined entirely by our desires?¹⁵ According to this view, perhaps some of my reasons are tied to objective values like health, but others are tied simply to my desires, such as my desire to watch as much soccer as possible. However, the general phenomenological point that I've been pressing holds across the board; in any given case, when one is faced with the sort of reflective pressures that I've been discussing, it appears as though an appeal to desire-independent values is what must ultimately justify a proposed course of action. A hybrid view of this sort will accordingly fail for the same reasons that DBR fails.

Before moving on to the next phase of the argument, I should address the role that desires play in this bottom-end theory of reasons. One might wonder if, when choosing between objects of relatively equal value (precisely how equal is an important but difficult question¹⁶), one's tastes and preferences can ultimately play a deciding factor in what one has reason to do. For example, assuming that the pursuit

¹⁵ Ruth Chang has recently offered such an account in her "Can Desires Provide Reasons for Action?"

¹⁶ Precisely what it means to speak of projects of "comparable value" is a very big question. Susan Wolf offers some interesting discussion of this question, although she too appears to wrestle with the question—without any final resolution—of how far to push the idea that (a) values can be grouped and lexically ordered (to at least some extent), and (b) values should be responded to on the basis of these orderings. On the one hand, she holds that "The view that one should love what is good in proportion to its goodness [which] evidently presupposes that some things are better than others" is an option that is "intelligible but horrible" (Wolf: 230). On the other hand, she suggests that some sort of "proportionality requirement" demanding that one's desires be proportional to the independent value of their objects is required by common sense (Wolf: 232). Wolf winds up with a "middle-ground" view that seems to allow for an extremely modest proportionality requirement, but her final view on this issue is unclear.

of a college major in either Economics or Philosophy will be validated by the existence of (roughly) equal value properties lying along either career path, perhaps the fact that I am drawn more toward the latter is sufficient to justify my pursuit of it.¹⁷ But can one consistently subscribe to VBR and grant that desires can play this sort of tie-breaking role?

While it appears as though desires can help break these sorts of ties, what is really doing the relevant work is the desire-independent value of subjective satisfaction, not the presence of desires as such. If I want to study philosophy more than economics, for example, then it will be the presence of desire-independent values associated with studying philosophy (such as the enjoyment I'll get out of it) which will break the tie. The upshot is that if we are to speak of a tie-breaking role of desires, it lies not in the desires themselves, but rather in the fact that particular instantiations of the desire-independent value of subjective satisfaction are determined by the desires of the particular agent (so desires are, in a sense, simply the mechanism through which we have access to many forms of subjective satisfaction).

4.4 Taking the Phenomenology of Deliberation Seriously

4.4.1 The Objection

At this stage I hope to have demonstrated that VBR does a better job than DBR of capturing the existence conditions for normative reasons for action. Sincere reflection reveals that our very concept of a normative reason is that of a pro-consideration grounded in desire-independent values.¹⁸ Since values are what

¹⁷ There are actually several possible ways of articulating this sort of “desires can break close ties” view. First, one could argue that the fact I simply like philosophy more than economics makes it the case that I *ought* major in the former (where the tie is broken by the desire as such; i.e. not because of the values that attach to desire-satisfaction). Second, one could argue that the balance of values makes it the case that I am *permitted* to choose either option (so I can choose philosophy simply because I like it, but I don't have to). I take my treatment of this issue to be a rejection of both of these options.

¹⁸ It has occurred to me that some may be troubled by the implication implicit in VBR that all normative reasons are grounded in values. Does VBR rule out the possibility that some of our reasons aren't

ultimately justify actions, VBR is what must ultimately cash out our platitude about the justificatory force of normative reasons.

At this point in the dialectic, however, an objection looms. As things stand, one could still object to VBR by arguing that we shouldn't take conclusions stemming from the deliberative point of view too seriously in the first place. Perhaps it can be established that VBR is validated from the deliberative point of view, but on the basis of this fact alone (the objection might continue) one cannot draw any interesting conclusions about what practical reasons are like. You might think that you are responding to desire-independent values, but in fact you are merely chasing after the projection of your desires. The mere fact that the object of your desires appears valuable to you is of little interest; the appearance of value is simply the manner in which our desires manifest themselves in our deliberations.

This line of objection constitutes a threat to my argument by suggesting that, while it might appear that my desires aren't playing an important role in grounding my practical reasons, my desires are nonetheless doing all of the interesting work. The work my desires are doing, however, is "behind the scenes," so to speak; what I take to be desire-independent values are nothing more than the projection of my desires, and so these desires are what ultimately "count in favor" of performing certain actions.

Note that this objector isn't challenging my conceptual analysis of normative reasons; he is claiming instead that our concept of a normative reason needs to be overhauled because of problems that a VBR analysis runs into. These problems could

value-based? What about deontological reasons? Does VBR "load the dice" in favor of a broadly consequentialist moral theory over a deontological one? (Jonathan Dancy has raised related worries about a different view of the value/reason relationship; see his "Should We Pass the Buck?"). This is indeed a big question, and one that I cannot hope to resolve here. My inclination would be to argue that, due to intelligibility constraints, we have trouble making sense of reason-claims that aren't ultimately tied to value-claims. The upshot is that deontological reasons—even if they aren't seen as tied to value in their local application—must be ultimately explicable in terms of values if we are to make sense of them. For an argument in favor of this view by a someone working within what is considered to be a broadly deontological tradition, see Barbara Herman's "Leaving Deontology Behind" in *The Practice of Moral Judgment*.

be cashed out in one of two different ways. First, the objector could be claiming that our concept of a normative reason needs to be overhauled because the sorts of desire-independent values that VBR posits *simply don't exist*. One might be inclined to offer this objection on the basis of a general skepticism about the metaphysics of value that one thinks VBR is committed to (perhaps of the sort that inspired J.L. Mackie's argument from queerness). This particular objector is arguing that our concept of a normative reason is non-referring in much the same way that the concept of a magic spell is non-referring. The accompanying suggestion is that if values cannot play the right sort of justificatory role, desires are the next best alternative. This objector might go on to claim that, while it is hard to see how "desire-independent values" could exist as part of the furniture of the natural world, there is nothing at all "queer" about desires and their objects.¹⁹

Alternatively, the objector might not be concerned about the "queerness" of desire-independent values (there are no ontological worries here). Perhaps he is worried instead that, since there is no widespread agreement about what sorts of things are actually of value, our very concept of a normative reason—as understood by the VBR proponent—is incoherent. If there is indeed widespread, intractable disagreement about whether desire-independent values are tied to (for example) God's will, or Absolute Spirit, or pleasurable states of consciousness, then we may wonder (at least according to certain theories of concept formation) whether there really is a concept here for us to speak of in the first place. While we certainly don't need perfect overlap in our shared understanding of a concept's extension for the concept to be coherent, it would seem that we need at least marginal overlap if we are to show that there is much of a concept here to begin with. This line of objection could

¹⁹ This general move is roughly parallel to the one made by some non-cognitivists about moral statements whereby (a) it is conceded that our moral discourse is at least partially descriptive rather than purely expressive, yet (b) since there are no moral facts of the sort we think there are, we should start using moral discourse purely to express rather than describe.

accordingly be summed up as follows. Since (a) our concept of a normative reason as cashed out by VBR lacks coherence, and (b) substituting in a DBR analysis would bring some semblance of unity and order back to the concept (by unifying diverse conceptions of values under the umbrella of “desire projection”), we can conclude that the projectivist’s DBR view of reasons makes for a worthwhile replacement concept.

While I find the first version of this objection more threatening than the second, in the sections that follow I will attempt to rebut them both at the same time. I will do this by showing that, not only is there significant overlap in our thinking about certain desire-independent values (such that just about everyone can agree that certain properties are valuable), it will also turn out that there is nothing mysterious about these properties (and hence no question about whether they actually exist).

Before taking up this challenge, and to drive home the point that more needs to be said to validate VBR, it is worth noting that there are a number of philosophers who are quick to agree that we take ourselves to be tracking values, but are nonetheless rather insistent that VBR (or some close relative of it) is false.²⁰ Clearly these philosophers don’t think that the validity of the phenomenological points raised earlier is sufficient to establish VBR. It is thus important that we avoid the mistake of assuming that a mere survey of these sorts of considerations is sufficient to prove VBR (despite the fact that some thinkers who argue for VBR seem to think that this sort of survey might suffice).²¹

²⁰ For example, despite denying the veracity of VBR, Michael Smith (with Phillip Petit) in “Backgrounding Desire”; Simon Blackburn in *Ruling Passions* (Blackburn: 253-4); and Stephen Darwall in “Because I Want It” (Darwall: 142), all seem to agree with the general Section 4.2 claim that we take ourselves to be tracking desire-independent values. Blackburn and Smith are on my interpretation both proponents of DBR. While Darwall does spend much time arguing for DBR, he is quite vocal in his rejection of VBR.

²¹ It is worth noting that while Dancy claims to be sensitive to this shortcoming (Dancy (1): 31), it isn’t entirely clear that he himself avoids it, at least with respect to one of his central lines of argument in favor of VBR (Dancy (1): 35). As suggested earlier, Scanlon could also be accused of making this mistake.

4.4.2 Framework For a Response.

To reply to our central objection, we first need to lay out some terminology. Of the concepts that we possess, many are meant to pick out particular things in the world. Our concept of health, for example, refers to properties of living things that are conducive to their longevity and vital functioning, whereas our concept of witchcraft refers to supernatural forces. Let us refer to concepts that purport to pick out objects that in fact don't exist "non-referring concepts." The concept of witchcraft, for example, would appear to be non-referring, on the assumption that the sorts of supernatural forces that it ostensibly refers to don't actually exist.

Next, let us define as a reductive analysis any analysis of a concept C that attempts to account for the role that C plays in our thinking by "reducing" its purported referent R1 to some alternative referent R2. You offer a reductive analysis of C, in other words, by reducing its purported referent—which is assumed to be either non-existent, dispensable, or mysterious in some way—to something that actually exists and is either indispensable or unmysterious. There are two basic goals one could have in offering a reductive analysis of this sort. First, one could intend to *eliminate* R1 by reducing it to R2, thus bypassing both the need and the meaningfulness of speaking of R1 at all. Second, one could intend to *vindicate* R1 by reducing it to R2, thus allowing us to go on speaking of R1 by accounting for it in terms of R2. One might be tempted to offer either form of reduction in response to different possible concerns about R1. I might think I should eliminate R1 through reduction either because I think R1 doesn't exist, or because I find R1 mysterious or tantalizingly dispensable. Alternatively, I might think that a certain reduction will vindicate R1 by demystifying it in some way. As I'll soon argue, however, one can only have good reason for reducing a given R1 in very specific circumstances.

Perhaps some illustrations of these two forms of reduction would help. First, we could offer a reductive-eliminative analysis of witchcraft by arguing that since its purported referent (supernatural properties) doesn't exist, we must understand processes that are supposedly explained by witchcraft as explicable in terms of natural properties (so we can account for poor crop yields by pointing to facts about overused and depleted soil rather than supernatural agency). In this case our goal is to eliminate the referent "supernatural properties," not vindicate it by explaining it in terms of something else. Alternatively, the behaviorist project of reducing mental descriptions to facts about behavior could be interpreted as a reductive-vindicating project, insofar as the goal is to allow us to go on using mental descriptions despite the fact that we might initially have seen them as either dispensable or mysterious or both.²²

Which sort of reductive analysis of normative reasons is our projectivist recommending? Put in terms of "desire-independent values," the projectivist's central reduction is clearly eliminative. He is, after all, claiming that there are no such things as desire-independent values (on either interpretation of the objection), and that we are accordingly better off reducing talk of desire-independent values to talk of desires. This isn't to say that the projectivist is out to eliminate talk of values altogether. Indeed, the projectivist's larger project might include a reductive-vindicating analysis of the general category of values (i.e. values that aren't stipulated to be desire-independent). This objector might accordingly claim that our very notion of a value will make much more sense if we start thinking of values in terms of the objects of our desires. Doing so will "vindicate" not only our talk of values (by demystifying them), but our talk of normative reasons for action as well. In any event, our focus in what

²² To give one more example, in *What We Owe To Each Other*, Scanlon's "buck-passing" account of value could be understood as a reductive-vindicating analysis.

follows will be on reductive-eliminative analyses, since this is the sort that is currently threatening VBR.

This brings us back to my discussion of the temptations one might have for offering a reductive-eliminative analysis and my warning about whether we are justified in acting on these temptations. I suggested that one might be tempted to offer such an analysis for a number of different reasons; one might, for example, think that a given concept should be eliminated because it is mysterious in some way or tantalizingly dispensable. Contrary to these temptations, however, let me propose the following principle: the only cases in which it is appropriate to offer a reductive-eliminative analysis of a concept are those in which the purported referent of the concept does not exist. We can think of this as the “if it ain’t broke, don’t try to fix it” principle of conceptual fidelity. Since a lengthy discussion of the usefulness of having stable concepts would take us too far afield, I will limit my justification for this principle to an appeal to common sense. If a concept (i) turns out to be reasonably well-formed, (ii) plays roughly the same role in the thinking of a large number of people, and (iii) clearly refers to real, existing properties, a reductive-eliminative analysis would be inappropriate.

It seems perfectly acceptable, for example, to offer a reductive-eliminative analysis of witchcraft, since the purported referent of this concept doesn’t exist. Alternatively, based on the assumption that our concept of health is *not* non-referring (i.e. there actually are properties of living things that are conducive to their longevity and proper vital functioning), it seems rather inappropriate to offer a reductive-eliminative analysis of our concept of health.

The upshot is the following. Based on my discussion in Section 4.3, we have seen that our concept of a normative reason is tightly connected to the concept of a desire-independent value (at least as I have defined this term). The objection stated at

the beginning of this section, however, claims that we should start thinking of reasons as tied to desires instead. This objection, in other words, amounts to a reductive-eliminative analysis of our concept of a normative reason for action. Based on the general principle that I've just been discussing, however, this sort of move would be warranted only if our concept of a normative reason for action is non-referring. But is this the case? Are there no desire-independent values to which our concept of a normative reason for action refers?

4.4.3. Our Concept of a Normative Reason Does Refer

To answer this question, we need to look more closely at the sorts of desire-independent values that we, upon reflection, treat as proper grounds for our practical reasons. My goal in the previous section was to demonstrate that we take the former to ground the latter, regardless of whether in any particular case the sorts of values that we think exist actually do. But can we say a bit more about the character of these values? What makes it the case that a certain thing has desire-independent value? Is it just a brute fact that some things have value and others don't?²³

The general answer I want to propose is the following. At least when we focus on a certain subset of value-claims, what makes it the case that certain things have desire-independent value is that they are conducive to making a human life go well (understood in a very pedestrian sense). When I speak of a human life going well, I will be speaking "in the abstract," in two different senses. First, I mean to say that, *ceteris paribus*, a given person's life will go better for him if he possesses a certain value property. This doesn't mean that we can't think up extravagant cases in which possessing a property thought to be generally conducive to human well-being will

²³ It is worth noting that most of the thinkers working in the VBR tradition either say little about what sorts of things have value, or make various value claims without saying much about *why* certain objects have value and others don't.

actually prove non-beneficial. The idea is rather that we can identify properties that are *pro tanto* conducive to human well-being, and that this conduciveness is what accounts for the value of these properties.²⁴

Second, the sorts of reason-claims that I have in mind are meant to be abstract in the further sense that we can posit these claims without worrying about whose well-being we are looking at. When I claim that we can all agree that V is a genuine value (because it is conducive to human well-being), I mean that (at a minimum) for any person A, we can all agree that from the point of view of A, it would be good for A to have V—and hence A has *pro tanto* reason to secure V for himself—regardless of whether we happen to be A, care about A, dislike A, etc.²⁵

Perhaps some examples would help. Think back to some of the value claims that were made in Section 4.3. At least part of what seemed to be reason-giving about my friendship with Linda involved her possession of certain qualities—like kindness and a good sense of humor—that seem to make life better for those at the receiving end of them (and perhaps for the one who possesses them as well). Similarly, part of what seemed to justify my pursuit of philosophy was the value inherent in the actualization of my intellectual capacities (a value that plausibly contributes to my well-being).

Perhaps my example of the perceived value of popularity serves as an even better illustration. Not only does it seem to me that what I thought was a value wasn't; part of my evidence for thinking this turns on the fact that, whereas my pursuit of popularity left me rather unhappy, looking back I see that people who didn't chase

²⁴ It is worth noting that even in cases in which one is made worse off by acquiring a *pro tanto* value (e.g. one's intellectual achievement leads to murder by a jealous competitor), the explanation that is ultimately given for how one is made worse off will itself involve reference to some other constituent of human well-being. The upshot is that, while the particulars of certain situations might turn a *pro tanto* good into an evil, the reason-givingness of considerations of well-being is still what determines where the relevant all-things-considered reasons lie.

²⁵ I am trying to remain agnostic on the issue of whether the sorts of values at issue are agent-neutral or agent-relative.

after this false value (people who opted instead for friendships based on “deeper” things) seemed to lead better lives than I did. The candidacy of popularity as a value, in other words, is undermined by the fact that it isn’t actually pro tanto conducive to a human life going well.

In citing these sorts of cases, I would guess that most readers will find the value candidates I have been discussing fairly plausible. However, it is once again irrelevant (at least at this stage) whether these purported values are real values (perhaps not all of them are as conducive to well-being as I think). The lesson is that what lends plausibility to the claim that these are genuine values is in each case an explanation that links them to facts about what makes a human life go well. These sorts of facts are, in other words, about as good a place for a justification to bottom out as one could ask for.

Note that this general view is compatible with a wide variety of specific views of what makes a human life go well. The idea that many values are ultimately tied to facts about what makes a human life go well—or “eudaimonistic facts,” as we may call them—is one that many people can agree with, ranging from Aristotle to thinkers like Nietzsche and Thrasymachus. To figure out what specifically constitutes human well-being might be no easy task; this is something that Aristotelians will simply have to argue over with Thrasymacheans by looking at other arguments, empirical evidence, etc. My claim is simply that all of these people are working with the same general concept of normative reasons for action—that is, of (a) pro-considerations tied to desire-independent values, (b) a large subset of which are grounded in what I’ve been calling eudaimonistic facts.²⁶

²⁶ To clarify, I am not claiming that our concept of a normative reason for action links all reasons to eudaimonistic facts. Many people, for example, will be equally satisfied by justifications for actions that link them to the survival of animals on a distant part of the planet. The claim is that there is enough agreement to conclude that there is at least one type of justification that will always satisfy us; namely one that bottoms out in eudaimonistic facts.

Could anyone deny this? I can imagine someone claiming that while our normative reasons are indeed tied to desire-independent values, the values in question are all tied to God's will, and hence have nothing to do with human well-being. The threat that this particular worry represents brings us back to the second interpretation of our projectivist objection. This was the view that if enough people disagree on even the most general categories of desire-independent values—such that there is little overlap even on the question of whether eudaimonistic facts represent real values—then the projectivist's central move will begin to look more plausible. The obvious line of reply is to demonstrate that the widespread disagreement posited does not exist, which in this case involves showing that even among many religious views, reason-claims grounded in eudaimonistic facts are found widely satisfying.

And this does indeed appear to be the case. Eudaimonistic facts are often given the right sort of currency by the faithful, whether they come by way of divine commands that one look out for the well-being of others (“love thy neighbor”), appeals to what are presented as reasonable concerns about how well things are going to go for one in the afterlife, and general claims to the effect that our God is a loving God, and one who cares about how well things are going for humans. We might even argue that there is a tacit limiting condition of sorts on the authority of many religious edicts, whereby a divine command is reason-giving only if it can be understood as not poisonous on a large scale to general human well-being. How exactly this limiting condition plays out in the religious story that is told is somewhat open; different people of faith who all agree that there is a limiting condition could account for it different ways.²⁷

²⁷ It is pressures of this sort that lead Robert Adams to build a limiting condition into his own divine command theory that restricts the reason-givingness of divine commands to those issued by a loving God. See (Adams: 264, 281).

It doesn't matter whether we are inclined to account for this limiting condition by arguing that (a) God wanted it this way (as things may seem to the religious), or (b) drafters of religious views recognized the independent normative weight of this limiting condition and built it into their religious views to ensure their legitimacy (as things may seem to the atheist). Demonstrating that a proposed action would promote human well-being seems sufficient to establish the *prima facie* plausibility of the action's associated reason-claim in a way that rival reason-claims calling for the widespread (and long term) destruction of human goods are unable to match.

4.4.4 The Answer To the Objection

We can now see that our concept of a normative reason for action is tied not only to desire-independent values, but moreover, (at least in many cases) to desire-independent values that are grounded in eudaimonistic facts. The question now is: has the discussion to this point given us good reason to think that our concept of a normative reason for action is *not* non-referring? While there may be disagreements over which specific properties are actually conducive to human well-being, I think we can safely say that there is enough overlap between rival conceptions of human well-being to warrant a positive response to this question.

Perhaps some examples would help. Take the following two properties: (1) the property of suffering from the most horribly excruciating and debilitating terminal illness imaginable, vs. (2) the property of being healthy. Is there really any deep uncertainty as to which of these two properties (if either) is more *pro tanto* conducive to human well-being? If the answer to even this one question is "no," then we have established that there is at least one property that answers to our concept of a normative reason for action (and more specifically, to the concept of a desire-independent value that underwrites such reasons). To link this claim up to the two

versions of our objection, not only is it the case that we seem to have widespread agreement in the view that being healthy is a genuine *pro tanto* value, there is also nothing mysterious about the property of being healthy.

While there is room to argue about how many properties of this sort there are, we can surely think of at least a few more of this sort. It seems to me, for example, that the following properties are all widely acceptable value candidates due to their conduciveness to human well-being: (1) having the ability to figure out what is of value and how to respond to it; (2) being free from states of absolute, unbridled terror or unending hopelessness; and (3) having the basic means for one's survival (such as food, shelter, etc.). I assume that not even a Thrasymachus or a Nietzsche could disagree with these claims.

In addition, there is another category of value candidates we could cite that, while falling short of universal acceptance, are plausible enough that we can confidently defend them against any non-fanatical objections. Items on this list might include: being engaged in meaningful social relationships, having a sense that one matters or has worth, being excellent at what one does, and having certain positive psychological experiences (such as pleasure or a sense of aesthetic appreciation). While it isn't hard to imagine some people—even of a relatively non-fanatical stripe—objecting to some of these items²⁸, I imagine that these objections will be held by a relatively small minority, and moreover, a minority that I am (perhaps optimistically) confident that we could prove wrong through extended discussion.

This last claim about resolving conflicting views about what goes on our list of value properties is an important one. It would be easy to think that since many people

²⁸ It has been pointed out to me that even religious views that require one to renounce pleasure or abandon any sense of self-worth don't contradict the view I have been defending. The sorts of religious views in question actually require that these things be seen as values, simply because the purported value of one's sacrifice of pleasure (for example) would come to little if pleasure wasn't itself valuable. Accordingly, even views that demand the renunciation of these goods can agree that these things are values; they simply think that other things are even more valuable.

disagree about what should go on the list of human goods (especially when we start getting into details), the number of values that we can legitimately put on the list must be very small. However, the mere fact of disagreement doesn't rule out the possibility that there are many more value properties out there than can sustain near-universal agreement (i.e. it may be the case that many people are simply wrong about what contributes to human well-being). As my own adolescent "value of popularity" case illustrates, in some instances even the people who held the wrong view can be brought to acknowledge that they were mistaken.

If, then, I have shown that our concept of a normative reason for action is *not* non-referring, then based on our general principle concerning the acceptability of eliminating referents, we must conclude that it would be inappropriate to offer the sort of reductive-eliminative analysis that the projectivist objector recommends. Since there do exist in the world at least some properties that answer to our deliberative questions about what to do (which are moreover properties of a fairly "non-mysterious" sort), it simply doesn't make sense to insist that we replace our concept of a normative reason with something else. If it ain't broke, don't try to fix it.

Perhaps an analogy would help show where our objector goes wrong. It is as if we were reflectively committed to cooking with only healthy ingredients, and our projectivist objector was arguing that since no one food is really any healthier than another (or that there is no such thing as "healthiness"), we should admit that we have purchased the items in our pantry simply because we like them (and so our likes, rather than considerations about conduciveness to health, are what really "count in favor" of purchasing broccoli rather than bacon). However, since we take ourselves to be shopping for only healthy foods, and since some foods are in fact healthier to eat than others, I don't see how we can avoid concluding that this projectivist is trying to do something inappropriate.

With this said, and if my arguments have been successful, I have not only shown that the phenomenology of deliberation recommends VBR over DBR, I have also shown that we can discard such considerations only at the expense of making an unnecessary act of violence against our concept of a normative reason for action.

4.5 Belief-Based vs. Fact-Based Normative Reasons

In addition to arguing for VBR and the notion that desires are not the central, underlying determinants of our reasons for action, I want to discuss a further distinction between two possible understandings of normative reasons and argue that beliefs don't play a central role in determining our reasons for action either. While this point will perhaps not be as central to the arguments that come in later chapters as VBR, it will have an important role to play. (The argument for this view is also parasitic on the argument VBR, in that it turns on appeals to the phenomenology of deliberation.)²⁹

The distinction I have in mind involves two different senses in which we talk about the reasons that an agent faces. The first sense in which one can “have a reason” to perform some action—which I will refer to as “belief-based reasons”—involves identifying the reasons that an agent has relative to the beliefs that she has at a given time (to include perhaps some assessment of the reasonableness of those beliefs). If it is time for me to go to the airport to pick up a friend, and I believe that my car keys are on my desk (since that is where I last put them), then I have a belief-based reason to head for my desk. The fact that someone moved my keys from my desk to the mantle without my knowledge will be in this case irrelevant; I still have belief-based reason to head to my desk because this is where my reasonable belief tells

²⁹ For a detailed discussion of the parallels between the first-person “aims” of beliefs and motivated desires, see Richard Moran’s discussion of beliefs, desires, and the transparency condition in *Authority and Estrangement*.

me to go. This sense of “having a reason” is the one we typically evoke in our assessments of behavior. So long as an agent’s beliefs are reasonable, we typically won’t hold it against her that the world did not in fact conform to those beliefs. This is also the sense of “having a reason” that we evoke when excusing an action that fails because our mistaken beliefs led us astray.

In contrast, we can also identify what I will refer to as “fact-based reasons.” This stronger sense in which one can be said to have a reason—which admittedly comes up less often in our conversations about reasons on a day-to-day basis— involves picking out the reasons that exist for us relative to the way that the world actually is. To return to our missing keys example, despite the fact that my mistaken belief about my keys’ whereabouts is reasonable, I can be said to have a fact-based reason to head for the mantle rather than for my desk. Indeed—and this is the important point—from my own point of view as a deliberating agent, fact-based reasons are what I am typically most concerned with. I need to pick up my friend from the airport; accordingly, what I have reason to do is to go to where my car keys actually are. While my mistaken belief will no doubt lead me to my desk, and while I will no doubt *excuse* my trip upstairs to my office by citing my mistaken belief, the success conditions of my action are, from my own point of view, established by the way the world is, and not by how my beliefs are.

In assessing the usefulness of the distinction that I have just drawn, it is important that we not read too much into our tendency to think of questions of “rationality” as questions of responsiveness to belief-based reasons. While it is true that we rarely (if ever) think of ourselves as being “irrational” for acting in a way commensurate with our reasonable, albeit mistaken beliefs, this is irrelevant to the point at issue. What is important is this: from the first-person deliberative point of view, our aim is to respond appropriately to fact-based reasons. Accordingly, the

normative reasons that we seek in our deliberations (in addition to conforming to VBR) will be fact-based. Call this the Fact-based Reasons Thesis (FBR).

Since our interaction with the world is always mediated by our beliefs, it is true that we will invariably execute this task of seeking out and responding to fact-based reasons by acting on the basis of our best understanding of how the world is. Accordingly, there is an important sense in which talk of fact-based reasons will have little bearing on our actual deliberations; I admit no useful practical guidance falls out FBR. But since my interest isn't to provide practical guidance, this objection doesn't speak against FBR. What we are trying to get at in our deliberations are fact-based reasons, not belief-based reasons, even if all we can ever guarantee is some responsiveness to the latter.³⁰

To give one final, intuition-priming argument for FBR, imagine for a moment that an all-knowing genie offers to give me advice before I set off to fetch my keys, and that I ask him what I have most reason to do. Suppose the genie replies by asking: "What do you mean 'most reason?' Would you like to know what you have reason to do given your reasonable albeit mistaken beliefs (i.e. do you want belief-based reasons), or would you like to know what you have reason to do given where your keys actually are (i.e. do you want fact-based reasons)?" Given my aims as a deliberating agent, it should be clear that what I most want is an answer to the second question. The fact-based reasons in any given case are what I am ultimately trying to track.

³⁰ To put this point in a way that recognizes the indispensable role that our beliefs play in our decisions, we could say that our actions will be successes (at least in part) to the extent that the belief-based reasons that we wind up pursuing are at the same time fact-based reasons (i.e. to the extent that our beliefs are true).