

The Objective Significance of the Subjective

1 Introduction

Sometimes it appears to us that we have (or lack) a good reason for some doxastic attitude. Epistemic rationality requires us to take that apparent reason into account. But do we gain a good reason to have the doxastic attitude just because it appears to us that we have a good reason? On a simple and tempting view inspired by analogies with ethics, the answer is ‘no’. The appearance of a reason in favor of some act or attitude might provide one with a good excuse for performing that act or forming that attitude. But excuses for responding in certain ways aren’t necessarily reasons in favor of responding in those ways. Even if we could be more rational or irrational for responding or failing to respond to appearances of reasons, we can distinguish between being rational and responding to good reasons in favor of attitudes and acts: rationality, on one popular view in ethics, provides a kind of excuse, not (necessarily) a kind of justification.

In this paper, I will explain how this picture might be resisted and how we might argue that the appearance of support gives rise to real support in epistemology. I will argue that this picture stands or falls with some deeper views about the structure of epistemic value. In particular, I will argue that if fundamental epistemic values like accuracy ought to be understood *non-teleologically* – so that what it is for accuracy to be epistemically valuable is for it to be *respected as a standard* rather than *produced as a goal* – we should think that appearances of reasons can generate real reasons in epistemology. I will not try to argue for a non-teleological picture

of epistemic value in this paper. I do that elsewhere.¹ I am more interested in showing that one defensible non-teleological picture of epistemic value undermines the idea that appearances of reasons don't necessarily give rise to real reasons (and merely give rise to excuses). Indeed, I think friends of the normativity of epistemic rationality *must* accept such a picture.

2 Divorce Pictures

Consider the following cases:

Faultlessly False Beliefs. Zane falsely but faultlessly believes that a certain potion would kill Jane when it is in fact a cure for all her ills. He recognizes that giving Jane this potion would be prohibited by decisive reasons if it really had the powers he takes it to have. But he gives her the potion anyway, and she drinks it. Her ills are cured, unbeknownst to him.

Misleading Appearances. Jill has strong but misleading evidence that flipping the switch would blow up Bill's house and that Bill is inside. She is in a position to know that she has this evidence, and recognizes that killing Bill would be disfavored by decisive reasons in the situation that her evidence suggests to obtain. But she flips the switch anyway. It actually dumps a pot of gold into his house, unbeknownst to her.

Something seems to be going seriously wrong in both cases. What is wrong isn't just that Zane and Jill have partly false beliefs and misleading evidence. While there is some difference between these cases, what's wrong seems similar: the agents are failing to act in the ways that the appearances suggest would be right or wrong. This failure seems like a real failure, not just the appearance of a failure.

¹See my "Veritism without Instrumentalism".

What is the best way to describe the failure? One thing on which most normative theorists would agree is that the agents in these cases are not as praiseworthy as they could be. Indeed, many would say that Zane and Jill are highly blameworthy. This would provide a partial answer to our question. Blameworthiness is certainly a flaw. But it is, at bottom, a flaw in the *agent*. It doesn't follow that there is some further flaw in the *act*. So, we have to ask: does the problem with the agent exhaust what's going wrong in these cases?

It is increasingly common for people to say 'yes'. Many normative theorists would say: "While Zane and Jill may be blameworthy, their acts are still supported by the reasons, and are actually right." They might also say that Zane and Jill have "subjective reasons" counting against their acts, and are for this reason failing to be fully rational. But lots would say this while admitting that subjective reasons do not provide real constraints on action. Subjective reasons aren't real reasons, on this picture. They are just apparent reasons.

Let's generalize this approach and give it a name. According to this approach, we can separate

- (1) what there are real normative reasons for some agent to do (or not do)

from

- (2) what this agent would be rational or irrational in doing (or not doing)

and from

- (3) what this agent would be praiseworthy or blameworthy for doing.

According to the view at issue, we can hold that

- (4) while the appearances concerning what there is reason for someone to do are relevant to (2) and (3), they are not relevant to (1).

Call (4) the *Practical Divorce View*.

One might think that a parallel view ought to be upheld in epistemology. One might insist that epistemologists, too, can separate

(1*) what there are real epistemic reasons for a subject to believe

from

(2*) what this subject would be epistemically rational or irrational in believing

and from

(3*) what this subject would be epistemically praiseworthy or blameworthy for believing.

According to the analogue of the Practical Divorce View,

(4*) while the appearances concerning what a subject has reason to believe can affect (2*) and (3*), they cannot affect (1*).

Call (4*) the *Epistemic Divorce View*. If this view were true, there would be weighty implications for many debates.

Consider the New Evil Demon Problem for reliabilist accounts of epistemic justification. Friends of the Epistemic Divorce View could say that while an honest reliabilist view will imply that the subject in the demon world has no genuine epistemic reason to believe what she believes, she is excusable and rational in believing what she believes. This is because it does appear, from her perspective, as if her experiences were reliable indicators of the worldly facts, and she forms her beliefs in line with the appearances. Those who reject reliabilism on the basis of the New Evil Demon Problem are conflating (1*) with (2*) and (3*).

This move would resolve some internalism/externalism debates in epistemology. One debate is between accessibilists and anti-accessibilists about justification. Simple accessibilists hold that

a factor F can bear on the justificational status (“J-status”) of some subject S’s belief only if S is in a position to recognize that this is so.

Since subjects in demon worlds are not in a position to recognize that their beliefs are unreliably formed, simple accessibilists deny that this fact bears on the J-status of their beliefs. This claim is tempting. After all, surely it would be a mistake to fault the believers here, and to call them less than fully rational! But not so fast: this exclamation conflates the question of whether a believer should be faulted with the question of whether her beliefs are in good standing. These are different questions. Since, in general, such questions can be divorced, this argument for accessibilism about justification gets nowhere.

The anti-accessibilist about justification can replace the earlier claim with the claim that

a factor F can bear on whether a subject S is open to criticism or praise for some belief B only if S is in a position to recognize that this is so,

or with the claim that

a factor F can bear on whether a subject is rational or irrational for holding some belief B only if S is in a position to recognize that this is so.

Accordingly, one can deny accessibilism about J-status while upholding it about the agent’s openness to epistemic criticism and praise, and about epistemic rationality and irrationality. In general, such properties are separable. We must recognize this in ethics, and so we must in epistemology too.

The Epistemic Divorce View also has implications for newer disputes about peer disagreement. If some epistemic peer disagrees with you but does so for what turn out to be bad reasons, our intuitions pull in different directions on the question of how you ought to react. On the one hand, because you only know that this person is your epistemic peer and is as likely to be right about questions of this type as you are, it would seem irrational for you to “stick to your guns” and ignore the fact that this peer has come to a different conclusion. On the other hand, from a God’s-eye point of view, your peer happens to be going wrong in this case, and you happen to be responding correctly to the evidence. Why should his mistake exert any real normative pressure on you or make you less justified? Surely you would do the right thing if you did stick to your guns, and you would be responding correctly to the real normative facts: your peer’s judgment adds nothing to those facts. That much can be justly said, at any rate, if we accept the Epistemic Divorce View.

All of this is just to say the Epistemic Divorce View matters. It is attractive partly because it matches what is a popular view about the relations between reasons, rationality and questions of criticism and praise in the practical sphere, and partly because it offers a neat way of reconciling conflicting intuitions about important cases and hence of partially resolving some important debates in epistemology.

3 Teleology as a Key Motivation for Divorce

The mere fact that the Epistemic Divorce View happens to be structurally similar to an initially attractive view in the practical sphere is not by itself a sufficient reason for endorsing it. There may be significant disanalogies between epistemology and ethics on this score. Moreover, the analogous view in ethics is not without its opponents. There are hard cases for the view that the ‘ought’ that matters in

deliberation is completely indifferent to perspectival factors.² And many have not been happy to concede to Kolodny that rationality doesn't give us real reasons. Intuitions to the contrary persist. There is a booming industry of attempts to vindicate the intuitions. This is so even though it is also clear that the Practical Divorce View would, if true, offer a neat way of reconciling conflicting intuitions.

So, just from what's been said so far, there is not obviously sufficient reason to embrace the Epistemic Divorce View. Can more be said?

Well, I think that there is a much deeper rationale for the Epistemic Divorce View. We can appreciate this rationale by considering a similar rationale for the Practical Divorce View.

One deeper rationale for the Practical Divorce View appeals to the link between reasons and value and some independent facts about value in the practical sphere. Facts about the distribution of value over outcomes are, in plenty of cases, perspective-independent. What matters for right action is, in large part, the distribution of value over possible outcomes. Even if we go nonconsequentialist, it would be crazy to deny this. In many cases, we don't seem to care in any nonderivative way about the distribution of merely expected value. Of course, from the deliberative point of view, the issues coincide: in trying to increase actual value, what we do is increase expected value. But when we reflect on the fact that the appearances may mislead us, we know that there will sometimes be reasons to regret what we do in the name of expected value.

A similar thought is tempting in epistemology. It is what can lead one to think that some form of reliabilism must be the right view about what makes for good evidence, and for justified belief rather than merely excusable belief. After all, reliabilists are in the unique position of being able to say: "Beliefs that are supported

²See Jackson (1991) and Zimmerman (2008).

by good evidence in our sense, and that are justified in our sense, are objectively likely to be true. Surely you care about that! But merely being rational in the narrow sense of living up to the appearances won't help you on that score. You could be perfectly rational and honor the appearances while being egregiously unreliable." There is something right here. Facts about the objective probability of getting it right given that one uses certain methods are perspective-independent. If it were true that these facts are what we ought to care about, we would have sufficient reason to think that some form of reliabilism is the only viable view about good evidence and justified belief.

These programmatic remarks can feel persuasive. But they conceal some crucial assumptions. The most crucial assumption is that value is teleologically structured in the epistemic domain in Scanlon (1998)'s sense: that what it is to be valuable is to be "to be promoted", and that the evaluative ranking of responses is fixed by the extent to which more valuable states of affairs and fewer disvaluable states of affairs would be promoted by these responses.

I agree that what we ought to do is correctly respond to value. I also agree that, in the practical case, some values are best understood teleologically. The value of cake plausibly consists in the fact that its nature gives us reasons to produce more instances of it, and the disvalue of filth plausibly consists in the fact that its nature gives us reasons to eliminate instances of it.

Like Anderson, Scanlon, and others, I am skeptical that *all* values in the practical sphere can be understood in this way. Some values ought most fundamentally to be honored, respected, or given some kind of recognition in our deliberation while producing more and more of them is optional albeit good. Other values ought fundamentally to be produced and respected while the reasons for production and the reasons for respect may draw on different features of the value (e.g., visual art).

Which fitting response is fundamental and explains the fittingness of others will, I think, vary from value to value and disvalue to disvalue. But I agree that there are plenty of values that are best understood teleologically.

But it is far from clear that a teleological picture is plausible for the fundamental epistemic values. Indeed, a teleological picture arguably fails for the epistemic value on which reliabilists tend to focus most: namely, accuracy. I've argued for this conclusion elsewhere. I won't repeat the argument here: what I'm interested in is how a non-teleological view could undermine the Epistemic Divorce View.

4 How a Non-Teleological View Could Undermine the Epistemic Divorce View

I will now explain why the arguments for the Epistemic Divorce View could fail if we rejected a teleological picture of epistemic value, and why, more strongly, we should reject the Epistemic Divorce View if we reject a teleological picture of epistemic value. I will draw on an analogy in the practical case: I will show how the Practical Divorce View can sometimes fail, and then extend the type of argument that arises in this special sort of practical case to the case of interest.

Let me first stress again that teleology in the sense at issue is a claim about value. It is a view on which what it is for something to be valuable is for its nature to provide us with reason to produce more instances of it, and on which what it is for something to be disvaluable is for its nature to provide us with reason to eliminate instances of it. If we adopt a buckpassing account on which what it is for anything to be valuable just is for there to be reason for us to have certain positive responses to it, it becomes unclear why a teleological picture should be the only game in town, or even the default view.³ Why couldn't the fitting response to some fundamental

³It's worth recalling that Scanlon (1998) tries to use the buckpassing account to undermine tele-

values be something other than promotion?

So, following Anderson, Scanlon and others, I would deny that teleology is a trivial conceptual truth about value. When we claim that persons are valuable, this does not trivially entail that we have reason to see to it that there are more people in the world. We are better understood as making a claim that entails that we have reason to respect people and take their interests at least as seriously in our thinking about what to do as our own. While we have strong reasons to make existing people sufficiently happy, we have no strong reasons to make more sufficiently happy new people. Relatedly, when we claim that friendship is valuable, it is not obvious that we are *merely* claiming that there is reason to produce more instances of friendship. It is better to see this claim as implying that we have special reasons to treat people who are our friends in certain ways and refrain from treating them in other ways.

When a value is best understood in a nonteleological way, appearances can give rise to real reasons. This is clearest when the response whose fittingness grounds the value is something like respect.

It is easiest to see this by considering examples like:

Ingenious Marital Reconnaissance. Jane worries that her husband Dale would cheat on her if he had the opportunity. She decides to test this hypothesis in a clever way. Using her extraordinary costuming skills, she manages to dress up to look exactly like a quite different woman on whom she suspects Dale would have an instant crush. In this disguise, she starts regularly showing

ology. As Hooker and Stratton-Lake (2006) note, there couldn't be a plausible deductive argument from the truth of the buckpassing account to the falsity of teleology. One could *consistently* claim that (i) for X to be valuable is for there to be reasons for us to have certain responses to X in our acts and attitudes, and (ii) that the only response whose support by reasons explains value is promotion. What I am saying is a more charitable reading of Scanlon (1998): the point is not that the truth of the buckpassing account entails the falsity of teleology, but rather that a large range of nonteleological views about the nature of value become clearly intelligible once the buckpassing account is endorsed. This puts a burden of proof on the long tradition of assuming that value is just by definition to be understood teleologically.

up outside his workplace to flirt with him. Dale believes on this misleading but compelling evidence that he is interacting with a different woman. He flirts with her in turn. Indeed, he makes plans to have a romantic evening with her when he next thinks that his wife is out of town, and to not tell his wife a word about it. That's precisely what he thinks he is doing next. All the while, he knows that his wife would be outraged if she found out. But he does not care: he thinks he is getting away with it, and has a great time.

Dale would be a fool if, when Jane claims later that he lacks sufficient respect for their relationship, he replied: "But isn't it clear that I have great respect for our relationship? It was you, after all, with whom I had that wonderful evening. How could I have made my respect for this relationship more apparent?" This speech would be absurd. Indeed, surely Dale would be speaking falsely when he claims that his acts manifested any respect for the relationship.

Our relationships have value for us and those involved, and part of this value plausibly consists in a nonteleological normative fact: the fact that we ought to have some kind of respect for these relationships. If we ignore what the appearances suggest would be a failure of respect and act in that way anyway, *we thereby exhibit a real failure of respect*. The failure is a failure that we have reason – real reason, connected with real value – to avoid. So, in these cases, the appearance of failure can constitute a genuine failure. The appearance of reasons not to act in certain ways can generate real reasons not to act in those ways. Call this kind of argument for rejecting the Practical Divorce View (in limited cases) the *Respect Argument*.

Sure, plenty of practical values aren't best understood in this non-teleological way. So the Respect Argument has limited force in the practical domain, and the Practical Divorce View may well apply in many cases. But it is, I think, an open question whether any fundamental epistemic values ought to be understood

teleologically. If no epistemic values ought to be so understood, but epistemic norms are in the end about being responsive to epistemic value, the Respect Argument may generalize to undermine the Epistemic Divorce View.

It is straightforward enough how it could generalize. One might embrace an epistemic Kantianism on which

what it is for accuracy to be (epistemically) valuable just is for us to have (epistemic) reason to respect accuracy as the norm of belief.

One might also think that

to respect accuracy as the norm of belief just is to be disposed to believe that P only if it *appears* that P is likely enough true.

Given these claims, one could advance the following revised Respect Argument:

1. Given epistemic Kantianism, accuracy's (epistemic) value consists in the fact that there is (epistemic) reason to respect accuracy as the norm of belief.
2. If it appears that you have sufficient reason to avoid forming some belief because it appears that there is too great a risk of inaccuracy, you would fail respect accuracy as the norm of belief in forming this belief *even if there is, objectively speaking, little or no risk of inaccuracy.*
3. So, if it appears that you have sufficient reason to avoid forming some belief because it appears that there is too great a risk of inaccuracy, then there is, in fact, sufficient reason to avoid forming this belief.

The core idea is simple. Suppose that Jones believes on the basis of the higher-order appearances that the first-order appearances on which he has been forming his beliefs are radically misleading. If he said to himself (even if only subconsciously),

“Who cares? I’m just going to keep taking the first-order appearances at face value, even though they are entirely misleading”, it would be impossible to regard him as respecting accuracy as the norm of belief. This intuition remains even if the higher-order appearances turned out to be misleading, and the process of forming a belief on the first-order appearances was completely reliable. So, if there is reason to respect accuracy as the norm of belief, there would be reason for Jones not to believe in accordance with the first-order appearances even if these appearances turned out to be perfectly trustworthy.

Accordingly, there is an argument from a nonteleological account of the value of accuracy to the falsity of the Epistemic Divorce View, just as there was an argument from a nonteleological account of certain practical values to the inapplicability of the Practical Divorce View in certain practical cases. Since epistemic Kantianism is one attractive nonteleological view about epistemic value and it is plausible that respecting a norm requires what appears likely to conform to the norm,⁴ we would have reason to reject the Epistemic Divorce View if we had reason to reject a teleological picture of epistemic value.

5 Limitations?

Let’s more carefully consider this question:

- Q. Is it *always* necessary to heed appearances of epistemic reasons in order to count as respecting accuracy as the norm of belief?

How we ought to answer Q will determine how badly wrong the Epistemic Divorce View might be. It is clear that it is sometimes necessary to heed appearances of epistemic reasons to count as respecting accuracy as the norm of belief. If it appears

⁴At least in some sense of ‘respect’. Cf. Darwall (1977)’s notion of recognition respect.

that some new belief-forming method would be grossly unreliable, one would clearly fail to respect the norm of accuracy if one adopted this method.

Not all cases are obviously like this. Suppose a reliable person invents a machine that she claims to be perfect at detecting whether people are in pain. This machine might have a perfect track record so far, and I might know this. Now she tries to use the machine on me. No dice: I am in intense pain, and her machine says I'm not. Given my prior knowledge about her machine, its verdict does yield apparently strong evidence by my own lights that I am not in pain. But I can dismiss this apparent evidence. If she tries to argue with me, it would be fine for me to reply, "I just *am* in pain!" It is sometimes OK to beg the question in that way. This is not to say that introspection is infallible. I *know* it is not. My introspective abilities might be less reliable in general than her machine. Still, I can dismiss the verdict her machine delivers. Even after entertaining self-doubt and recognizing the general facts about the quality of the belief-forming method on which I know I'm relying, I can permissibly beg the question.

But this isn't a victory for the Epistemic Divorce View. It does not suggest that we should revise our judgments about the earlier cases. Introspection of pain is an unusual case. It affords one of the few cases of a belief-forming method that can be understood *radically* rather than liberally or conservatively.⁵ There is no reason why we cannot be radicals about some methods, liberals about others, and conservatives about yet others. Only in cases where we ought to be radicals will it

⁵I here borrow some terminology from Pryor (2004: 353-4), but modify and extend it in a way suggested by Brian Weatherson. In the usage here, a method M is to be understood conservatively when one needs antecedent justification for believing that M is reliable in order to gain justified beliefs via M. M is to be understood liberally when one doesn't need antecedent justification for believing that M is reliable in order to gain justified beliefs via M, but when there can be evidence which is sufficient to defeat the support that M provides for a belief. M is to be understood radically when one doesn't need antecedent justification for believing that M is reliable in order to gain justified beliefs via M, and where there could be no evidence that would be sufficient to defeat the support that M provides for a belief in any case.

be plausible that one can ignore higher-order appearance without failing to respect the norm of accuracy.

Barring some general argument for radicalism, I see no reason to think that this case reflects a common pattern. Indeed, in cases where a method is best understood conservatively (e.g., extended cognition) or liberally (e.g., testimony), higher-order appearances that seem to cast doubt on the reliability of the method seem just as significant as evidence that rebuts the verdicts of the method. One would fail just as much to respect the norm of accuracy if one ignored the higher-order appearances as one would if one ignored rebutting first-order evidence.

We are left with a mixed verdict. This is not unsettling, since we are left with a similar verdict in structurally analogous circumstances in the practical case.

Consider Huck Finn, to see the analogy. He believed that in caring for Jim, he was doing the wrong thing. He took the racist norms of his community to be authoritative. Yet he could not bring himself to do otherwise. We take him to be praiseworthy for this, even though he does seem to be failing to comply with enkratic norms (e.g., “If you believe you oughtn’t A, then don’t A”). We don’t take him to be praiseworthy just for having his act produced by a “morally reliable” character trait: at some important level, he is motivated by the right reasons, and does properly care for moral value. This is a case in which it is possible to act against the appearances of reasons while retaining respect for the value that grounds these reasons. But not all cases in the practical arena can be understood like this. In the two cases with which we began, we certainly don’t want to regard the agents as having appropriate respect for the people at issue, and part of the reason for this is precisely that they are acting against what the appearances suggest would be morally right.

I’ll leave it open exactly how we ought to explain the fact that appearances of reasons are sometimes relevant to respect for value, and sometimes not. The key

point is that responding to appearances of epistemic reasons is generally necessary for respecting accuracy as the norm of belief. If so, it follows from epistemic Kantianism that appearances of epistemic reasons can generate real epistemic reasons.

6 Implications

Let's consider the upshots. What would follow if we rejected the Epistemic Divorce Picture on the basis of a non-teleological picture of epistemic value?

One of the most striking implications is that epistemologists could not separate the hypological (i.e., agent-oriented matters like blameworthiness and praiseworthiness) from the deontic (i.e., act/attitude-oriented matters like support by reasons, obligation and justification) as clearly as ethicists can.

Since it is clear that some practical values are indeed to be understood purely teleologically, it should not be surprising that there could be reasons for an agent to act in some way even when the appearances suggest that acting in that way would be entirely wrong. It can seem plausible that giving Jane the life-saving potion in *Faultlessly False Beliefs* was the right thing to do even though Zane nonculpably believed that, by giving her the potion, he would kill her. This is because what really matters in this kind of case is avoiding actual harm to Jane: the importance of this matter completely trumps whatever agent-relative reason there might have been that flowed from the fact that Zane disrespects Jane's humanity here.

But when a respect-based understanding of some value is plausible, it will no longer be possible to separate the subject's openness to criticism or praise from the wrongness or rightness of her attitudes. Norms of respect ground agent-oriented criticism and praise. If you fail to respect some value that ought to be respected, you are *ipso facto* open to criticism; similarly, if you respect some value that ought to be respected, you are *ipso facto* blameless on the score of this obligation, and indeed

typically praiseworthy.⁶ Norms of respect and disrespect cannot be blamelessly violated. Such norms just are the basis for praise and criticism. You can't truly disrespect a value that ought to be respected without automatically being criticizable to some degree; if you aren't criticizable, that is actually a reason for thinking that, in some way or other, you didn't truly disrespect the value.

What we would see in epistemology is that what Scanlon (2008) calls the deliberative and critical uses of a principle are not independent. Recall that, as Scanlon has it, the deliberative use of a principle decides what an agent or subject has most reason to do (in a broad sense), and the critical use decides how the agent deserves to be assessed for taking the reasons stated by the principle into account. When the value that grounds the principle is a value that ought primarily to be given recognition or respect and whose maximization is optional, the fact that the principle is to be used critically in some way on some occasion partly determines the way in which it is to be used deliberatively on that same occasion. What the agent or subject has most reason to do (in the broad sense) will partly be determined by how good a job the agent does in making certain considerations her motivating reasons. While this is by no means always or even often the case with norms (particularly not with those stemming from values that are primarily to be promoted), there are some crucial norms for which it will be the case. If we reject a teleological picture of fundamental epistemic value, and opt for some respect-based nonteleological view, fundamental epistemic norms will coincide in their critical and deliberative uses.

⁶You might try to generate a counterexample to this claim by imagining cases in which someone is blameless or not fully praiseworthy because she is subject to some kind of autonomy-precluding manipulation or coercion. But this would have to rest on a mistake.

It is only necessary for respecting a norm in the sense at issue that one autonomously heed the appearances regarding whether the conditions of application of the norm obtain. If someone failed to heed such appearances but only because she was manipulated or coerced, it would be implausible that she failed to respect the relevant norm. She may have had great respect for it: the manipulation just prevented this respect from manifesting in what she ended up believing or doing.

The conclusion also has implications for some disputes in normative epistemology. When it comes to issues of peer disagreement, we would have to accept something more like Kelly's Total Evidence View. Even if you happen, as a matter of fact, to be more reliably responding to first-order evidence and your opponent is uncharacteristically confused, you will often be obliged to back off a bit unless his confusion is clear to you; the only clear exceptions will be when your belief is formed by a method that can rightly be understood in radical terms (e.g., introspection).

We would also have to reject radical kinds of externalism about epistemic justification favored by people who think that hypothetical and deontic considerations in epistemology can be fully orthogonal. It is a mistake to think that the subject in the demon world who believes on bases that he blamelessly – indeed praiseworthily – takes to be perfectly good has nothing by way of genuine support for his beliefs. This subject isn't falling short in the fundamental duty to respect accuracy as the norm for belief. Indeed, he's doing a good job on that score. It may well be the only score that matters, if accuracy is the sole intrinsic epistemic value, inaccuracy is the sole intrinsic epistemic disvalue, and both values are understood nonteleologically. But whatever we say about monism vs. pluralism about intrinsic epistemic value, there's a strong consideration that counts in favor of our subject's doxastic attitudes, and that gives him something by way of real support.

A related implication is that epistemologists cannot separate rationality and support by normative reasons as sharply as ethicists plausibly can. We should follow Scanlon, Parfit and others in thinking that rationality is, at its core, a matter of responding to apparent reasons. Apparent reasons are believed reasons or reasons that appear to obtain in some non-doxastic sense. Part of rationality so understood is wide scope. To be fully rational, we must avoid conflicting states of the following forms, where 'A' picks out any commitment-involving attitude like intention or

belief:

- S believes that she has decisive reasons not to A yet As.
- S believes that she lacks sufficient reasons for A-ing yet As.
- S believes that she has decisive reasons to A yet fails to A

Another part of rationality so understood is narrow scope. On one plausible view, it consists in complying with process requirements like the following:

- If the non-doxastic appearances suggest that S has decisive reasons not to A, rationality requires S not to A.
- If the non-doxastic appearances suggest that S lacks sufficient reasons for A-ing, rationality requires S not to A.
- If the non-doxastic appearances suggest that S has decisive reasons to A, rationality requires S to A.

It is easy enough to see how these requirements could fail to be reason-implying. We can run Kolodny-style anti-bootstrapping arguments to make this conclusion plausible in many cases in the practical arena. It is not true that you get any strong reason to A simply because the appearances suggest that A-ing is in accord with the reasons. After all, if the potion you've got is poisonous, what you really ought to do is not give it to your friend. You may be blameless for giving it to her if your evidence suggests that this potion is not poison. But what you in fact have reason to do is provided by the real facts about the potion and the circumstances. This is plausibly because the values at issue here are best understood teleologically: well-being is primarily to be promoted, and ill-being is primarily to be prevented.

Of course, not all examples in the practical case will provide plausible bases for anti-bootstrapping arguments. As we saw in *Ingenious Marital Reconnaissance*, it is not plausible that what Dale did really was OK. He failed to have the right kind of respect for his relationship because of the way in which he acted in the light of the appearances. The asymmetry between this case and the earlier case owes to the differing character of the values at issue: while the value of well-being and the disvalue of ill-being are rightly understood teleologically, the value of Dale's relationship with his wife is best understood in a nonteleological, respect-based way.

When the values that provide the reasons are to be understood in this way, appearances of reasons will give rise to real reasons. And if fundamental epistemic values like accuracy are best understood in a respect-based fashion, there may be an asymmetry between the significance of rationality in epistemology and in ethics.

References

- Bach, Kent. 1985. "A Rationale for Reliabilism." *The Monist* 68: 246 – 63.
- Broome, John. 1999. "Normative Requirements." *Ratio* 11: 398 - 419.
- Dancy, Jonathan. 1992. "Externalism for Internalists." *Philosophical Issues* 2: 93 - 114.
- Engel, Mylan. 1992. "Personal and Doxastic Justification in Epistemology." *Philosophical Studies* 67: 133 – 50.
- Firth, Roderick. 1981. "Epistemic Merit, Intrinsic and Instrumental." *Proceedings and Addresses of the American Philosophical Association* 55: 5 – 23.
- Fumerton, Richard. 2001. "Epistemic Justification and Normativity" in Steup, M. (ed.) *Knowledge, Truth and Duty*. Oxford: Oxford University Press.
- Hooker, Brad and Stratton-Lake, Phillip. 2006. "Scanlon versus Moore on Goodness" in Horgan, T. and Timmons, M. (eds.) *Metaethics after Moore*. Oxford:

Oxford University Press.

Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101.3: 461 – 482.

Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114: 509 - 63.

Kolodny, Niko. 2008. "How Does Coherence Matter?" *Proceedings of the Aristotelian Society* 107.1: 229 – 263.

Korsgaard, Christine. 1997. "The Normativity of Instrumental Reason" in Cullity, G. and Gaut, Berys (eds.) *Ethics and Practical Reason*. Oxford: Clarendon Press.

Littlejohn, Clayton. 2009. "The Externalist's Demon." *Canadian Journal of Philosophy* 39.3: 399 – 434.

Parfit, Derek. 2001. "Rationality and Reasons" in Dan Egonsson, Björn Pettersson and Toni Rønnow-Rasmussen (eds.) *Exploring Practical Philosophy*. Aldershot: Ashgate.

Parfit, Derek. 2012. *On What Matters*. Oxford: Oxford University Press.

Pryor, James. 2004. "What's Wrong with Moore's Argument?" *Philosophical Issues* 14: 349 – 378.

Raz, Joseph. 2005. "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1.1.

Ross, Jacob. Ms. "Rationality, Normativity and Commitment".

Ryan, Sharon. 2003. "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies* 114: 47 – 79.

Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.

Scanlon, T. M. 2007. "Structural Irrationality" in Brennan, G., Goodin, R., Jackson, F. and Smith, M. (eds.) *Common Minds: Essays in Honor of Philip Pettit*. Oxford: Oxford University Press.

Scanlon, T. M. 2008. *Permissibility, Meaning and Blame*. Cambridge: Harvard University Press.

Schroeder, Mark. 2007. "Teleology, Agent-Relative Value and 'Good'." *Ethics* 117: 265 – 295.

Sosa, Ernest. 2001. "For the Love of Truth?" in Fairweather, Abrol and Zagzebski, Linda (eds.) *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford: Oxford University Press.

Thomson, Judith Jarvis 1986. "Imposing Risks" in *Rights, Restitution and Risk*. Cambridge: Harvard University Press.

Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge: Harvard University Press.

Thomson, Judith Jarvis. 2008. *Normativity*. Chicago: Open Court.

Weatherson, Brian. 2008. "Deontology and Descartes' Demon." *The Journal of Philosophy* 105: 540 – 569.

Wedgwood, Ralph. 2002. "The Aim of Belief." *Philosophical Perspectives* 16: 267 - 297.

Zimmerman, Michael. 1997. "A Plea for Accuses." *American Philosophical Quarterly* 34.2: 229 – 243.

Zimmerman, Michael. 2008. *Living with Uncertainty*. Cambridge: Cambridge University Press.