

Below is a draft of mine from 2012 circulated and archived by email, containing ideas which eventually made it into 'An Epistemic Non-Consequentialism'.

See <http://www.kurtlsylvan.com/b-sides-demos-false-starts.html> for other old drafts.

From Teleology to Rationality's Insignificance (and Back)

This chapter has three aims. The first is to argue that epistemic rationality cannot have robust normative significance if one accepts any *prima facie* plausible version of *teleology about epistemic norms*. That is the view (roughly put) that whether certain epistemic norms are right depends on whether complying with them would constitute, or objectively likely constitute, promotion of fundamental epistemic value. Since the consequences of denying that epistemic rationality has robust normative significance conflict with orthodox internalism *and* externalism, we should be interested in non-teleological views. The second aim is to argue against teleology about epistemic *norms* by arguing against a deeper form of teleology about epistemic *value*. The latter leads, I will show, to swamping problems about the value of epistemic justification (and rationality). We *can* solve these problems by rejecting the view that what it is to be epistemically valuable is to be a goal that we ought epistemically to promote. We *cannot* solve some of these problems otherwise, contrary to what some optimists believe. Once teleology about epistemic *value* is rejected, epistemic teleology about norms is clearly false or unmotivated. The final aim will be to show how certain non-teleological theories I call *fitting response theories* can explain why epistemic rationality has robust normative significance.

With these aims in mind, here is how I'll proceed. In §1, I introduce several varieties of teleology about epistemic norms after explaining where this view and its competitors fall in the structure of epistemological theorizing. I also introduce teleology about epistemic value. I'll show how these views are dialectically connected, and why teleology about epistemic norms plausibly stands or falls with teleology about epistemic value. In §2, I explain why teleology about epistemic norms leads to pessimism about the robust normative significance of epistemic rationality. In §3, I will present my argument against teleology about epistemic value, by showing how it leads to some swamping problems. Then, I will show how the falsity of this deeper view undermines teleology about epistemic norms, by showing it to be either false or unmotivated. In §4, I discuss in greater detail a class of non-teleological theories that promise to solve these swamping problems and vindicate the normativity of epistemic rationality. In §5, I prepare for the next chapter by considering what the best fitting response theory would have to look like to *fully* vindicate the normativity of epistemic rationality.

1. The Underexplored Bedrock of First-Order Epistemology

1.1. *General first-order criterial theories vs. fundamental theories*

Epistemologists are often interested in nonnormative criteria for various normative properties or relations, in a broad sense of 'normative'. They want to know what necessarily *makes for* the presence of these properties or relations, in general nonnormative terms. Take noninferential epistemic justification, for instance. Different epistemologists offer different ways of filling in the blank:

Necessarily, S is noninferentially justified in believing that P iff _____.

The blank is to be filled in with a conjunction of general nonnormative conditions, such as *S's belief that P was produced by an unconditionally reliable belief-independent belief-forming process....*

Debates can then be had about the intensional adequacy of different views. Is it necessarily true that every subject who satisfies the conditions used to fill in the blank is noninferentially justified? Conditions can be added or subtracted depending on predictions about possible cases.

Call theories of this kind *general first-order criterial theories*.¹ Such theories are general, because they are not just listing conditions that can *in some worlds* make for the target normative status, but rather ones supposed to be necessary and sufficient for it in any world. They are also first-order. They are trying to answer questions about normative properties and relations are fixed by *nonnormative* properties and relations. They are *not* trying to answer questions about which facts *among all the normative facts* are the most fundamental, or to analyze certain normative facts in terms of *other sorts of normative facts*. Finally, they are *criterial* because they are just in the business of specifying general nonnormative criteria for the presence of various normative properties and relations in any possible world. When such theories are offered, their proponents do not typically try to explain in deeper terms why these criteria are the *right* criteria—not at great length, at any rate. As the game is normally played, it is enough to note that the criteria yield more plausible intensional implications than any competing criteria, that the criteria are simple, natural, and unified, and so on.

First-order criterial theorizing is not the only task that should interest epistemologists *qua* normative theorists. Even if we knew how to fill in a blank like the one above with general criteria, we can still ask *why* in deeper terms this is the right way of filling it in. Suppose for argument's sake that process reliabilism about noninferential justification is a true general first-order criterial theory. We can still ask: *why* is process reliability relevant in this way to epistemic justification?

Call a theory that will enable us to answer questions of this variety a *fundamental* theory. Epistemic teleology is one fundamental theory. The *veritist* version of this theory identifies the state of believing truly as the fundamental epistemic value and the state of falsely believing as the fundamental epistemic disvalue. On one articulation, it assumes that what it is for true belief to be epistemically valuable is for it to be a *goal* that we ought epistemically to promote by believing truly, and it then exploits a link between epistemic rightness and epistemic value to derive an account of epistemic rightness. If the rightness at issue is the rightness of certain *rules*, one version of veritist epistemic teleology might tell us that the right rules are those which, when followed, produce a high ratio of true beliefs to false beliefs in the long run. In a moment, I will discuss the details and varieties of epistemic teleology at greater length—right now I am simply mentioning it for the sake of illustration.

Some theorists (e.g., Alvin Goldman and William Alston) commit to fundamental theories and plausibly observe that these theories help to ground their criterial theories.² Veritist teleology fits well

1 The distinction to follow is modeled on Kagan (1992)'s distinction between “factor theories” and “foundational theories”. I've made the first term more exact, and changed the second term to avoid confusing epistemologists.

2 See, e.g., Goldman (1980: 32) and Alston (2005). The case of Goldman is nuanced. On the one hand, one finds teleological assertions like the following in the just-cited article: “The choice of a [doxastic decision procedure] clearly depends on the goals of cognition, or doxastic-attitude-formation.” On the other hand, one gets the impression elsewhere that Goldman would embrace process reliabilism as a true first-order general criterial theory regardless of what the fundamental theory is. He is happy to argue for it on the basis of its extensional superiority to other views, its naturalistic character, its elegance, and so on. Indeed, Goldman once stressed to me when discussing epistemic teleology that although he has sometimes asserted the view, he didn't mean to assert it *loudly*. Process reliabilism is more important to him. And it was first intended as a first-order criterial theory:

I want a set of substantive conditions that specify when a belief is justified. Compare the moral term 'right'. This might be defined in other ethical terms or phrases, a task appropriate to meta-ethics. The task of normative ethics, by contrast, is to state substantive conditions for the rightness of actions. Normative ethics tries to specify non-ethical conditions that determine when an action is right. A familiar example is act-utilitarianism, which says an action is right if and only if it produces, or would produce, at least as much net happiness as any alternative open to the agent. These necessary and sufficient conditions clearly

with reliabilism, even if it is unclear whether it is *needed* to motivate reliabilism.³ But fundamental theories are rarely subject to much discussion or dispute. The focus has been on criterial theories; many assume that it is possible to decide between such theories on the basis of familiar theoretical virtues like extensional adequacy, simplicity, etc. With no pressing need to decide on a fundamental theory, epistemologists are content to indicate their sympathies briefly and pursue less deep arguments.

Nevertheless, many epistemologists indicate commitment to epistemic teleology but neither defend this commitment nor explain how it meshes with the general first-order criterial theories they recommend. Comparably fundamental alternatives to epistemic teleology have scarcely been considered. In many cases this is striking, since it is unclear how certain general first-order criterial theories *could be* motivated by epistemic teleology. Consider how Bonjour expressed sympathies for veritist teleology in his coherentist years.⁴ It is notoriously hard to see how coherentist justification could serve the putative goal of maximizing true beliefs and minimizing false beliefs, barring a coherentist theory of truth. If anything, veritist epistemic teleology seems *undermine* coherentism.

1.2. *Why there must be more discussion about the correct fundamental theory*

This is one illustration of why there must be more discussion about fundamental theories. As I will argue, veritist epistemic teleology guarantees that there will be a significant gap between epistemic rationality and justification. It guarantees, indeed, that the problem of normative significance for rationality introduced in the last chapter is insoluble. People who deny a gap between rationality and justification or believe that the problem of normative significance is soluble should accept a different fundamental theory. Some commonly held views must go. Since I will argue that epistemic teleology is what we should reject, we must think about what will replace it. This is something few have done.

Some might say that we can engage in general criterial theorizing without taking a stance on the correct fundamental theory. This would be true if intuition together with other theoretical virtues were sufficient to discriminate between competing theories. But as I argued in earlier chapters, crucial intuitive disagreements are at a standstill in the literature on epistemic justification. Some key intuitions—like the one behind the new evil demon problem—can be taken seriously *only if* the distinction between justification and rationality is blurred in epistemology in a way in which it is no

involve no ethical notions. Analogously, I want a theory of justified belief to specify in non-epistemic terms when a belief is justified. This is not the only theory of justifiedness one might seek, but it is one important kind of theory and the kind sought here. (Goldman (1979: 1))

This combination of sentiments is visible in the Introduction to his (2012), where he discusses Selim Berker's case against epistemic teleology. For this sort of reason, my target in what follows is not *directly* Goldman or other reliabilists, except perhaps Alston. One might be a reliabilist and reject epistemic teleology. I do think this is not an easy combination of positions to defend. But it is not, I'll argue in §3.1.4, an incoherent combination.

3 When it comes to knowledge, I believe we should accept some kind of reliabilism (ideally virtue-theoretic) *regardless* of the correct foundational normative theory. As we saw in earlier chapters, however, I think that knowledge does not partially *consist in* justified belief. Knowledge is at least as basic as justification. As we've also seen, I do not think that our views about knowledge should strongly constrain our views about justification, and agree with Foley (2004) here. Knowledge without justification is a serious possibility, as is justification without (any) knowledge. Even if knowledge is *conceptually* more basic than justification, the specific *way* in which it is more basic may enable us to accept reliabilism about the former without accepting it about the latter.

4 Cf., for instance, Bonjour (1985: 7-8): "The basic role of justification is as a *means* to truth, a more directly attainable mediating link between our subjective starting point and our objective goal. We cannot, in most cases at least, bring it about directly that our beliefs are true, but we can presumably bring it about directly...that they are epistemically justified. And, *if our standards of epistemic justification are appropriately chosen*, bringing it about that our beliefs are epistemically justified will also tend to bring it about, in perhaps even longer run and with the usual slippage and uncertainty which our finitude mandates, that they are true."

longer blurred in the practical reason literature. As we will see in this chapter, this blurring is acceptable only if epistemic teleology is rejected. So, extensional intuitions about justification alone are not going to be sufficient to resolve the remaining internalism/externalism disputes. We need to get into deeper territory, and the territory partly involves what the correct fundamental theory will be.

1.3. *The varieties and structure of epistemic teleology*

Epistemic teleology is the orthodox fundamental theory. All versions of epistemic teleology view *conduciveness to fundamental epistemic value* as the feature needed to explain the correctness of proposed criteria for different epistemically normative properties and relations. There are, of course, many epistemically normative properties and relations. So one can envisage many different specific versions of epistemic teleology. As we will see, however, these different views are dialectically related. It is hard to embrace some without embracing other, deeper versions.

Let's focus first on a class of theses worth putting under the heading of *teleology about epistemic norms*. In the present use, 'norms' refers to any normative factors in the broad sense that *ought to guide what we do*, in a broad sense of 'do'. Norms in this sense contrast with *mere standards of appraisal*, which may not give rise to any relevant analogue of 'ought' implies 'can', and which are thus not normative in a *narrow* sense. Norms here include facts about normative reasons, obligations, permissions, etc. They do not *directly* include facts about what would be *best* or satisfy standards of *appraisal*, though such facts are *indirectly* relevant to what norms ought to guide us *given* further conditions—e.g., that it is *open to us* to make certain things have these good-making features.

Since we have been focusing on epistemic *reasons*, it is useful to start with the following more specific version of teleology about epistemic norms, which is reminiscent of Alston (1988, 2005)'s externalist account of the adequacy of “grounds” for belief:

(TELEOLOGY—GOOD REASONS) R is a good *pro tanto* epistemic reason to form a doxastic attitude D with respect to a proposition P iff R's truth makes it objectively more likely that one will promote basic epistemic value(s) in virtue of D(P)-ing than that one will promote basic disvalue(s).⁵

While it may be endorsed in isolation (as in Alston), this thesis isn't best viewed as free-floating. It is best viewed as following from the conjunction of a deeper teleological claim and a platitude about the relation between *rightness* and *reasons*:

(RIGHTNESS CONSTRAINT—REASONS) R is a good *pro tanto* epistemic reason to form a doxastic attitude D with respect to a proposition P at *t* iff R's truth at *t* would make it objectively more likely that D-ing with respect to P would be epistemically right D-ing than that it would be epistemically wrong D-ing.

(TELEOLOGY—DOXASTIC RIGHTNESS) Forming doxastic attitude D with respect to P epistemically rightly is a matter of in fact promoting basic epistemic values in virtue of D(P)-ing (and not basic disvalues).

The notion of *rightness* at issue here is an *objective* sense—the sense associated with evaluations of *correctness*, like “true beliefs are correct”. The concept of rightness or correctness at issue remains a

5 The *prima facie* plausibility of this view and other versions of teleology about epistemic norms will depend crucially on how 'in virtue of' is understood. As I argue in the Appendix, Selim Berker's recent attacks on epistemic teleology target an implausible, alien version of the view because they understand 'in virtue of' in a *causal* and *diachronic* sense, which is not charitably attributed to most epistemic teleologists. I say 'in virtue of' here ought to be understood in a *constitutive, synchronic, proposition-relative* sense: one would promote basic epistemic values *in virtue of D(P)-ing at t* if D(P)-ing would itself *constitute* the promotion of these values at *t*.

normative concept, even if it is true that all and only true beliefs are correct.⁶

The first claim should not be a controversial claim. It is plausibly a platitude—an instance of a general conceptual truth about the relationship between reasons and rightness. How could some reasons for believing be truly *good*, if the presence of them did nothing to raise the chance that one would be believing rightly or correctly? I see no plausible answer to this question. So given that *everyone* should embrace RIGHTNESS CONSTRAINT—REASONS, TELEOLOGY—GOOD REASONS would only be plausible *if* TELEOLOGY—DOXASTIC RIGHTNESS were accepted. So although TELEOLOGY—DOXASTIC RIGHTNESS and TELEOLOGY—GOOD REASONS are distinct versions of teleology, the latter stands or falls with the former.

Now, TELEOLOGY—DOXASTIC RIGHTNESS is *itself* not a self-standing thesis. Like TELEOLOGY—GOOD REASONS, it is best viewed as an implication of two deeper claims: a platitude about the relationship between rightness and value and an even deeper teleological claim about value. If one of these claims fails, TELEOLOGY—DOXASTIC RIGHTNESS will be threatened. The deeper claims are:

(VALUE ALIGNMENT CONSTRAINT—RIGHTNESS) Acting or forming attitudes rightly is a matter of performing the best acts or forming the best attitudes that are open to one.

(TELEOLOGY—VALUE ALIGNMENT) The best actions and attitudes are those in virtue of which one most efficiently promotes basic values.

For our purposes, TELEOLOGY—DOXASTIC RIGHTNESS is more helpfully evaluated in the context of *specifically epistemic* versions of these general theses—viz.:

(EPISTEMIC VALUE CONSTRAINT—EPISTEMIC RIGHTNESS) D(P)-ing epistemically rightly is a matter of forming the epistemically best attitude available with respect to P.

(TELEOLOGY—EPISTEMIC VALUE ALIGNMENT) Forming the best doxastic attitude D with respect to P is a matter of in fact most efficiently promoting basic epistemic values in virtue of D(P)-ing.


The fact that TELEOLOGY—DOXASTIC RIGHTNESS is a consequence of these deeper theses would probably not be made explicit even if people were explicit about their commitment to it. For many still assume without argument that value itself is to be understood teleologically, almost as a matter of definition. But as Scanlon (1998), Anderson (1993) and Parfit (2011) argue, this thesis is hardly trivial. *Promoting* a value is just one response among many possible responses. Promotion aside, values can serve as objects of commitment, they can be respected, esteemed, honored, protected, preserved, and so on. TELEOLOGY—VALUE ALIGNMENT implicitly singles out promotion *qua* goal as the only fundamentally correct response to basic value. This, we will see, is a deeply problematic assumption, and from its falsity we can argue forcefully against other versions of epistemic teleology.

Now, it is unclear whether anyone has ever envisaged *arguing* for TELEOLOGY—DOXASTIC RIGHTNESS on the basis of these claims. But dialectically speaking, this is irrelevant. What matters is (i) that the first claim is extremely plausible, since it is an instance of a general platitude, and (ii) that TELEOLOGY—DOXASTIC RIGHTNESS would be *unmotivated* if the deeper thesis TELEOLOGY—VALUE ALIGNMENT were false. Indeed, my strategy in the next chapter will be to argue against TELEOLOGY—DOXASTIC RIGHTNESS by refuting TELEOLOGY—VALUE ALIGNMENT and showing that the negation of TELEOLOGY—DOXASTIC RIGHTNESS *follows* from a much better *non-teleological* account of epistemic value alignment.

This helps us to understand the structure of epistemic teleology. While epistemic teleology is

⁶ For helpful discussions of rightness or correctness in this sense, see Wedgwood (2002a, 2003, 2007, Ms). Wedgwood himself is committed to epistemic teleology, and analyzes correctness in a teleological manner. Setting aside this disagreement, the notion of correctness or rightness I have in mind is like the notion he has in mind.

usually *explicitly* held only about certain items of epistemic interest like *reasons* and *justification*, these explicit theses *stand or fall with more basic teleological theses*. With the case of reasons in mind, we can usefully chart the dependencies as follows, where the upward pointing arrow means 'supports':



One consequence of this way of picturing the substructure of commonly endorsed teleological claims is that the status of teleology is independent of the questions about whether the *good* is prior to the *right*. For it should be controversial whether the good is itself to be understood teleologically! So even if the good were prior to the right, that would show nothing about the status of *any* teleological theses unless a teleological thesis were already assumed—namely, TELEOLOGY—VALUE ALIGNMENT. This is a crucial point. It shows us that the attractiveness of the thought that the good is prior to the right does nothing to motivate any version of teleology.⁷ It only does so if teleology about value itself is assumed.

2. From Teleology about Epistemic Norms to the Impotence of Epistemic Rationality

Rationality, as I argued in Chapter 2, consists in responding appropriately to *apparent* reasons. As I argued in Chapter 3, the relevant apparent reasons satisfy the following minimal internalist constraint:

Supervenience Internalism: If two subjects S and S* are non-factive mental duplicates at *t*, they have all the same apparent reasons at *t*.

Whether any stronger kind of internalism holds is doubtful, as I also argued in Chapter 3. But given this minimal internalist constraint, it is easy to argue that epistemic rationality is normatively impotent as such if any minimally plausible version of teleology about epistemic norms is true.

Before turning to that argument, some clarifications are in order. I will focus initially on *veritist* teleology, on which true belief is the basic epistemic value. Admittedly, if being epistemically rational is one of the *basic* or *underived* epistemic values, teleology about epistemic norms would *not* yield skepticism about the robust normative significance of epistemic rationality. Being epistemically rational would then trivially guarantee fulfillment of a basic epistemic goal. But it is exceedingly implausible that if there are basic epistemic goals at all, epistemic rationality will be among them. There is near universal agreement about this. Epistemic rationality clearly seems to have *derivative* significance. Of course, whether *derivative* significance is *instrumental* significance of any kind is, as we'll see, precisely the question that should concern us once we see the implications of veritist epistemic teleology. It is, I will argue, a great mistake to think that epistemic rationality has only instrumental epistemic value. But it will *not follow* even on my view that it has *underived* epistemic value. The conclusion I'll push is that there are other forms of epistemic value *derivation* beyond *instrumental* epistemic value derivation. Something can have *non-instrumental* but nevertheless *derivative* epistemic value (as I've also argued in Sylvan (forthcoming), *contra* Pritchard (2011)'s assumptions). Since, for the *teleologist*, derived epistemic value *is* instrumental epistemic value, the

⁷ This point is familiar from Scanlon (1998).

teleologist cannot plausibly hold that epistemic rationality is a basic epistemic goal. Otherwise it would have *underived* epistemic value for her. And it simply doesn't, even on my view.

Of course, there are other ways to reject veritism than to add epistemic rationality to the stock of basic epistemic values. One could add *knowledge* or *understanding* to the stock of basic epistemic values. At the end of this section, I'll consider whether doing so helps. It will not, as we'll see, help. So, while I accept veritism and aim to *vindicate* it by accepting a *non-teleological* account of epistemic value (as in Sylvan (forthcoming)), this commitment of mine is not indispensable.

2.1. *From Veritist Teleology to the Normative Impotence of Epistemic Rationality*

Here is a simple argument from veritist teleology about epistemic norms for the claim that rationality lacks robust normative import as such:

1. Rationality consists in complying with apparent reasons.
2. Given veritism, complying with apparent reasons by forming a doxastic attitude D(P) doesn't in all or even many possible cases make it objectively more likely that one will fulfill basic epistemic goals in virtue of D(P)-ing. Indeed, in many possible cases (e.g., demon worlds), it will *systematically* raise the objective chance that one will *not* promote basic epistemic goals in virtue of D(P)-ing, and indeed that one will *thwart* these goals.
3. TELEOLOGY—GOOD REASONS is true: R is a good *pro tanto* epistemic reason to form a doxastic attitude D with respect to P iff R's truth makes it objectively more likely that one will promote basic epistemic values in virtue of D(P)-ing than that one will promote basic epistemic disvalues.
4. If (2) and (3), complying with apparent reasons isn't as such complying with good reasons, given veritism. In many possible cases, complying with apparent reasons is complying with *bad* reasons.
5. So, complying with apparent reasons isn't as such complying with good reasons, given veritism; indeed, in many possible cases, complying with apparent reasons is complying with bad reasons.
6. If (5), rationality lacks robust normative import as such, given veritism.
7. So, rationality lacks robust normative import as such, given veritism.

There aren't many places where this argument could go wrong. (1) was defended at length in earlier chapters. (3) we are simply granting for the moment, since the broader strategy in this chapter will be a *reductio* on (3). (4) seems obviously true. After all, if (i) complying with apparent reasons does nothing *as such* to help one to realize epistemic goals in D(P)-ing and even may *thwart* these goals, and (ii) reasons are only good to the extent that complying with them *would* help to realize these goals, then apparent reasons are not necessarily good reasons, and may be bad. (5) and (7) are consequences of earlier steps. Finally, (6) is analytic, given how I'm using the terms. Like Kolodny, Broome, and others, I here just *mean*, by 'X has robust normative import', that X as such provides genuinely good reasons. (I would distinguish, of course, between *pro tanto* and decisive forms of normative potency.)

Where, then, could the argument go astray? (2) perhaps calls for some defense. So, one task of this section will be to remind the reader of some reasons to accept (2). I will do this in §2.4.

2.2. *Subjective Teleology: Does it Pose a Dialectical Problem?*

Before doing this, I want to forestall a couple of dialectical complaints. These rather than direct objections to the argument strike me as a bit more pressing.

Here is the first. Someone might grant the soundness of the argument, but complain that it shows that I have wrongly defined epistemic teleology. *Subjective* teleologists are likely to lodge this complaint. My reply will be that subjective teleology either concedes too much to me or simply changes the subject. The same problems about robust normative significance confront the subjective teleologist in a verbally different form. Subjective teleology *at best* fails to take these problems seriously. Since Chapter 4 was dedicated to explaining why we should, we can dismiss this refusal.

The *modest* subjective teleologist accepts the following dualist view about reasons and oughts: there are *subjective* and *objective* reasons, as well as *subjective* and *objective* oughts. He will then add that different notions of probability are relevant to stating teleology about these different sorts of reasons and oughts. Perhaps the claim will be that a reason for D(P)-ing is *subjectively good* that reason's obtaining make it more *subjectively* or *appearance-relatively* probable that one would promote basic epistemic values *in virtue of* D(P)-ing than that one would promote basic epistemic disvalues *via* D(P)-ing. A reason for D(P)-ing is *objectively good* iff that reason's obtaining make it more *objective* probable that one would promote basic epistemic values in virtue of D(P)-ing than that one would promote basic epistemic disvalues in virtue of D(P)-ing.

The modest subjective teleologist will then note that premise (3) could be read in two ways, since there are two versions of TELEOLOGY—GOOD REASONS: a version about *subjectively* good reasons, and a version about *objectively* good reasons. If it is understood as a view about subjectively good reasons, then premise (4) is false. For apparent reasons plausibly *are* subjectively good reasons. If understood as a view about objectively good reasons, the teleologist will simply concede the argument, but claim that the conclusion of the argument is not worrying.

To repeat the claims of the last chapter, this response involves so many concessions that its concluding strand of optimism is implausible. For it is hard to see how this response amounts to more than claiming that we ought rationally to be rational, and that we probably ought not always to be rational in any further sense of 'ought'. Talk of the “subjective ought” seems to be thinly concealed talk of the ought of rationality. Since what we subjectively ought to do plausibly consists in what we have sufficient apparent reasons to do, the subjective ought coincides with the ought of rationality.

Perhaps the objector thinks that there is not a serious question of the form: “Why be rational?” If so, my reply would be that there is, and that its answer isn't a trivial one like: “We should be rational because that is the way to be rational.” Of course, the objector's claim isn't *superficially* a tautology, since she uses talk of the subjective ought in her explanation. But we can simply raise all the same questions about the subjective ought that we can about the ought of rationality. These questions are not empty questions, to which the answer is: “Obviously, the subjective ought *doesn't* have robust normative significance as such.” That they are not is made evident by the last chapter, in which we saw that this position conflicts (*inter alia*) with most epistemologists' beliefs about the defeat of epistemic justification (which I've argued is conceptually distinct from rationality).

If the objector *isn't* making this move, her response concedes my argument without undermining its dialectical force. We want an explanation of why we really ought to be rational, or why rationality necessarily has *some* robust normative force, even if it is not decisive. On the present horn of the dilemma, the objector agrees. Yet to concede that there is no explanation and agree that rationality does not have any robust normative force as such *if* epistemic teleology about the *objective ought* concedes my argument. It does not undermine the importance of the argument. For, once again, we saw in the last chapter that conceding this conclusion is in tension with what most epistemologists

—internalists and externalists alike—believe. If the truth of the conclusion made no difference to epistemology, it wouldn't matter. But it does make a difference. So it matters.

Modest subjective teleology, then, can only briefly change the subject. If one distinguishes between subjective and objective oughts (or reasons), the question on the table then simply becomes: why do subjective oughts (or reasons) matter? Given the conclusions of the last chapter, they had *better* matter, since otherwise most epistemologists' beliefs about (e.g.) the defeat of justification would be false. If the modest subjective teleologist concedes that they don't matter *by* conceding that they do not, as such, have an impact on what we objectively ought epistemically to “do” or have objective reason to do, then the conclusion I want stands. If she does not mean to make this concession, she must add to her response. I do not see what addition could be made.

I've responded so far to the *modest* subjective teleologist, who is a dualist about oughts and reasons. A less modest subjective teleologist may claim that subjective oughts and reasons are the only oughts and reasons there are. This teleologist has a more direct response to the argument—namely, that (3) is the wrong version of teleology and that, once replaced, (4) becomes false. While this response is more direct, it also rests on unstable ground. Its defender would need to give us an error theory about apparently objective uses of concepts like REASON-FOR, EVIDENCE, and OUGHT.

Moreover, this objector would presumably not deny that there is a notion of epistemic *goodness* that is *not* perspective-dependent. This would be gratuitously error-theoretic. But admitting even this puts pressure on the immodest subjectivist. Sure, facts about epistemic goodness in this sense may not on their own give rise to epistemic oughts (or reasons), given the truth of some suitable relative of 'ought' implies 'can'; we cannot be obligated to do what we cannot do. But it is hard to see why the relevant relative of 'ought' implies 'can' would imply that only *apparent* epistemic goodness should be relevant to assessing how good our reasons are. This is not a general truth. To the contrary, it seems in general that a reason for A-ing is good only if complying with that reason makes a real difference to whether one's behavior would in fact be better, in the perspective-*independent* sense of 'better'. Certainly, the reasons for which we A must be reasons that can guide us—this is the core of the thought behind 'ought' implies 'can'. Still, there is better and worse guidance, and facts about *this* do not seem to be perspective-dependent in the way that the facts that *are* reasons might be perspective-dependent.

Let's take stock. The modest subjective teleologist cannot sap the dialectical force of the argument. At best, she can briefly change the subject, only to have the same questions arise with the same force about subjective oughts and reasons. The immodest teleologist can do better. But her position is too immodest. Its error-theoretic commitments cry out for motivation that hasn't been given.

2.3. *Are the Constraints on Robustness of Normative Force Too Strong?*

I turn to another indirect, dialectical objection. Remember that I followed Kolodny, Broome, and others in insisting that a source of requirements has *pro tanto* robust normative force only if it gives us genuinely good reasons as such. Like Kolodny and Broome again, I would add that a source of requirements has *outright* robust normative force only if these reasons are sufficient or decisive, so that we really *ought* to comply with these requirements. I am content to stipulate these senses for “pro tanto robust normative force” and “outright robust normative force”. The objector, I take it, would then insist that there are other ways in which rationality might have normative significance as such.

Of course, Kolodny, Broome, and others agree with this sentiment. Broome agrees that in the sense in which requirements of etiquette, grammar, etc., might be called “normative”, rationality is also normative.⁸ More strongly, Kolodny explicitly grants that rationality might have what he calls “evaluative” rather than normative (or “deontic”) significance, and indeed part of his error theory for

⁸ Broome opens his (2007a, b) and (2008) with this sort of remark.

explaining why rationality *seems* robustly normative involves an appeal to this point.⁹ In the last chapter, I similarly agreed that rationality might stand to having or responding to genuinely normative reasons as *excuse* stands to justification, in senses familiar from Austin (1956) onward.

If these are the *only* senses in which rationality has normative significance, then the conclusion of my argument remains interesting. It would be *surprising* if rationality coincided *only* with something like excusability as opposed to justification. Indeed, some epistemologists are keen on insisting that rationality *does* have this kind of significance. They partly assume this when they routinely conflate rationality with justification, as we saw in earlier chapters, and treat conditions for one as conditions for the other. But they are also often explicit about this. As Pryor (2001: 117) wrote: “It doesn't merely seem to me that the brain in a vat can form beliefs in a way that is epistemically blameless. It also seems to be the case that he can form beliefs in a way that is epistemically *proper*, and that the beliefs he so forms would be fully justified—despite the fact that they're reliably false.”

Beyond the potential evaluative significance of rationality and its relevance to questions of epistemic blame, what other kind of significance might it have? No one has proposed clear further alternatives, except perhaps to add—as Kolodny (2005: 512) anticipated and used as part of his error theory—that we ought *subjectively* to be rational. As we saw in the last section, this doesn't help. Barring further proposals, then, I think we can rest content with dismissing this objection. While the argument does only show that epistemic teleology precludes rationality from having robust normative significance in the senses initially discussed, this does not show that the argument is insignificant.

2.4. *Premise (2)*

The only other way the argument could go astray is with respect to premise (2). For it is the only premise that doesn't follow directly from points made in earlier chapters, from earlier premises, or from stipulations. Nevertheless, I cannot see how one could reject this premise.

This is obvious if apparent reasons satisfy Supervenience Internalism, which I defended in Chapters 2 and 3. Rationality is shared equally between the good case and the bad case. Make the case as bad as you like in a given possible world, so that responding to apparent reasons would ensure that one would reliably form false beliefs. If that were so, it couldn't be true in that world that complying with apparent reasons would make it objectively more likely that one would promote basic epistemic value, according to veritism. It could make exactly the reverse objectively certain.

Indeed, teleology strengthens the case for the claim, discussed in the last chapter, that mere appearances can be just as misleading with respect to the normative facts as beliefs can be. For whether the appearances are or are not misleading with respect to the real normative facts will turn entirely on *contingent* facts about the actual reliability of forming beliefs on the basis of apparent reasons. If they can be just as misleading, and we would *not* regard belief-relative narrow scope requirements as tracking the genuine normative facts when one's beliefs are false, we would have no resources within teleology for privileging the non-doxastic appearances. At that point, we can give something like Kolodny's bootstrapping argument. Rationality can, in many possible cases, require us to do what we lack sufficient genuine reasons to do, simply because the appearances are defective. That very fact deprives it of outright robust normative force in all cases, and—if the appearances are *really* defective—even of *pro tanto* normative force.

The problem here is not local to non-actual worlds. There are many cases in the actual world in which one has misleading higher-order “evidence”—that is, cases where it appears that certain actually good reasons are bad, or where certain actually bad reasons appear good. In these cases, there will be parallel problems to the ones that arise in thoroughly bad cases like demon worlds. These problems

⁹ See, e.g., Kolodny (2005: 571 and 554).

won't be *pervasive* in the way they are in demon worlds. But that is not comforting. If a view systematically makes the wrong predictions in a proper subclass of cases in a world, the view is still mistaken. It is just a much a mark against the normative significance of rationality *as such* if it requires us to do what we lack sufficient genuine reasons to do in a proper subclass of actual cases.

2.5. *Would It Help to Expand the List of Basic Epistemic Values?*

The argument stands. It is, of course, only an argument that *veritist* teleology about epistemic reasons undermines the normative significance of epistemic rationality. For it to be clear that the problem is a problem for *teleology*, parallel problems must confront other plausible versions of teleology.

I say other “plausible” versions because I rule out the idea that epistemic rationality is one of the *basic* epistemic values. This is in line with what most who assume epistemic teleology believe. It is also my own view. Of course, by the end of this chapter, we will see that we should not conflate the claim that rationality has *non-basic* epistemic value with the claim that it has *instrumental* value. This conflation is one of the core mistakes of teleology about epistemic value. Values—epistemic and non-epistemic alike—can fail to be *instrumentally* derived from other values without thereby having *underived* or *basic* epistemic value. Rationality, I will ultimately claim, has non-instrumental but nonetheless non-basic value. An advantage of drawing this distinction and rejecting teleology about epistemic value is that it allows us to secure the claim that epistemic rationality has only non-basic epistemic value with the claim that it is not, as such, instrumentally conducive to basic epistemic value.

Let's turn, then, to some more plausible non-veritist views. When epistemic value theorists propose to expand the list of basic values, the two most common candidates are knowledge and understanding. I'll consider these in turn. In both cases, I will argue that adding these candidates to the stock of basic values does not help to avert the kind of argument I offered in §2.1.

Take knowledge first. Knowledge is *factive*. For this reason, the second half of the second premise in the argument from §2.1 extends to show that in a range of key cases, adding knowledge to the stock of basic values will not help. After all, recall that complying with requirements of rationality in demon worlds will systematically guarantee that one will believe falsely. For this reason, it will systematically guarantee that one will not know. Accordingly, if one is looking for an *instrumental* explanation of why rationality matters—the *only* kind of explanation a teleologist could have—the explanation won't be improved by adding knowledge as a basic epistemic value.

What about cases in which responding to apparent reasons would not systematically guarantee that one would believe falsely, but rather just *fail* to increase the chance that one would believe truly? As far as I can see, the most that might be invoked by the knowledge-adding teleologist is the following. Responding to apparent reasons in believing P is necessary for knowing P. So, it does raise the chance that one will promote the basic epistemic value of knowledge in virtue of complying with these reasons—at least in cases where complying with these reasons would also not lower the chance that one would believe truly. Since that is what the defender of TELEOLOGY—GOOD REASONS needs to vindicate the robust normative significance of rationality in the cases at issue, she can do so.

Life is not so easy. Note that the response just floated on behalf of the teleologist requires:

(P) If X is a necessary condition for Y and X's obtaining does not also lower the chance that other necessary conditions for Y obtain, then X's obtaining raises the chance that Y's obtains.

This principle is false. In one way this is obvious. Suppose we're in a case where the other necessary conditions for Y besides X are *independently* guaranteed not to obtain. Here X's obtaining would not itself lower the chance that Y would obtain. Nevertheless, it clearly isn't true that X's obtaining raises

the chance that Y obtains. The chance is zero regardless of whether X obtains.

This holds for Gettier cases. By its nature, a Gettier case is one in which the objective chance that one would know if one believed rationally (and truly) is zero. Nevertheless, the fact that one believes rationally is not itself responsible for this maximally low chance. That is an independently ensured fact about the case. While it may be true that one must believe truly *because* one believes rationally to avoid being Gettiered, it is not *also* true that one's belief is Gettiered because one believes rationally in every Gettier case. If not, the antecedent of (P) may be satisfied in a Gettier case. Yet the consequent will not be true. The fact that a certain belief would be rational would not raise the chance this belief would be knowledge, though it may not be responsible for lowering that chance.

If all this is right, the suggestion on behalf of the teleologist fails. Indeed, the suggestion is bizarre, given the cases we're imagining. The only explanation the teleologist can give for why apparent reasons are always good is that they raise the chance that one would fulfill basic epistemic goals if one complied with them. Yet if knowledge is the basic epistemic value of interest, this strategy can't work for Gettier cases. In Gettier cases, a belief's being rational doesn't raise the chance that it would be knowledge. Yet rationality should have the same *kind* of significance as such whatever the case may be, if it has any significance. Knowledge-first epistemic teleologists cannot make sense of this apparent fact. For them, whatever normative significance rationality has is wholly contingent, and can vary from case to case in ways that seem bizarrely arbitrary.

So an analogue of premise (2) in our original argument would hold, even if the epistemic teleologist expands her list of basic epistemic values to include knowledge. The rest of the argument could then be slightly modified to yield the same kind of conclusion. So, adding knowledge to the stock of fundamental epistemic values won't help.

What about the other primary candidate one often sees—namely, understanding? This option strikes me as even less helpful. Remember: what we need is for the teleologist to offer a way to explain why apparent reasons for believing P are robustly normative reasons for believing P. Given TELEOLOGY—GOOD REASONS, there is only one way this could work: if complying with apparent reasons in believing P increased the chance that one would be fulfilling basic epistemic values in believing P. To evaluate the current suggestion, let us substitute in understanding as the relevant basic epistemic value. Then the proposal would seem to be this: R is a good epistemic reason for believing P only if R's truth would raise the chance that one would gain understanding in believing P.

This cannot be a helpful proposal. It is implausible on its face that good reasons for belief are considerations whose truth would increase the chance that we would gain greater understanding if we believed. There are *sufficient* reasons for believing many claims that wouldn't advance our understanding in the slightest. The fact that I see a speck of grain on the table is a sufficient reason given the rest of my condition to believe that there is a speck of grain on the table.

Maybe there is a trivial sense of “understanding” in which rationally believing P suffices for greater understanding. But even if that were true, there would be a different problem: it is not plausible that understanding of *this* sort is a basic epistemic value. No standard reasons for adding understanding to the list of basic epistemic values is a reason for adding “understanding” in this trivial sense to the list. Indeed, one standard motivation for adding understanding to the list would be a decisive reason *not* to add “understanding” in this sense to the list. The point of invoking “trivial truths” cases is that these are truths that wouldn't contribute to our understanding of anything.

So, there's a dilemma. Either understanding is understood (i) in a robust sense or (ii) in a trivial sense. If (i), the relevant instance of TELEOLOGY—GOOD REASONS is implausible. If (ii), understanding isn't a basic epistemic value after all. If so, the strategy we're considering can't work.

For this sort of reason, it is fair to conclude that teleology as such leads to skepticism about the robust normative significance of rationality. The problem is crystal-clear with veritism, once the

dialectical complaints considered in §2.2 and §2.3 are addressed. Adding *plausibly* to the stock of basic epistemic values doesn't help, as we've now seen. So, the same kind of argument can be used to show that these more pluralistic forms of epistemic teleology face the problem.

If this is right, teleology about epistemic norms is at the core of the problem of the normative significance of epistemic rationality. Of course, one might suspect that there are *other* routes to this problem. But rejecting epistemic teleology about epistemic norms leads to a natural solution to it.

3. Against Teleology

To get to this, I turn to an indirect argument against epistemic teleology about epistemic norms. I will argue against teleology about epistemic norms *by* arguing against teleology about epistemic value. My argument deliberately resembles arguments from Anderson (1993) and Scanlon (1998) in the practical sphere. They maintained that consequentialism is undermined once we appreciate that being valuable is not being “to be promoted”. They did so while agreeing that our reasons for action go hand in hand with facts about value, so that if all the evaluative facts are fixed, so are all the normative facts. This argument differs crucially from traditional deontologists' objections. For traditional deontologists, the normative facts are not fixed by facts about value. Scanlon and Anderson can reject that deontological thesis and agree that reasons and value are closely interlinked. They simply deny that value is itself what consequentialists presuppose it to be.

The rest of this section will amount to a defense of the first, second and fourth premises in the following argument, which I'll call the *Scanlonian Argument*:

1. TELEOLOGY—VALUE ALIGNMENT is false.
2. If (1), TELEOLOGY—DOXASTIC RIGHTNESS is also false.
3. So, TELEOLOGY—DOXASTIC RIGHTNESS is false.
4. If TELEOLOGY—DOXASTIC RIGHTNESS is false, then TELEOLOGY—GOOD REASONS is false.
5. So, TELEOLOGY—GOOD REASONS is false.

Let's turn to the arguments for (1), (2) and (4). The arguments for (2) and (4) are both simple, and were foreshadowed as I discussed the structure of teleology. So I'll spend most of my time on (1).

3.1. *Against TELEOLOGY—VALUE ALIGNMENT*

My argument against TELEOLOGY—VALUE ALIGNMENT is simple, though defending its premises will consume much of this section:

- i. TELEOLOGY—VALUE ALIGNMENT is true only if derived epistemic goodness is necessarily instrumental epistemic goodness.
- ii. But if derived epistemic goodness is necessarily instrumental goodness, the epistemic value of justification (and rationality) will be swamped by the more fundamental epistemic values from which it derives value.

- iii. Yet the epistemic value of justification (and rationality) is not swamped by the more fundamental epistemic values from which they derive value.
- iv. So, derived epistemic goodness is not necessarily instrumental epistemic goodness. (Indeed, the best way to solve these swamping problems is to recognize non-instrumental but nonetheless derivative epistemic value.)
- v. So, TELEOLOGY—VALUE ALIGNMENT is false.

Let me turn to defend the premises of this argument. Since (iv) and (v) are consequences of earlier steps, the only steps that could need defending are (i), (ii) and (iii). (iii), we will see, is obviously true, but its truth will become clear in the context of the discussion of (ii).

3.1.1. *Premise (i)*

First, a schematic defense of premise (i). Suppose the consequent of (i) is false: there is such a thing as epistemic value that is derived, but not derived along instrumental lines. TELEOLOGY—VALUE ALIGNMENT tells us that the best actions and attitudes are those which most efficiently promote—i.e., instrumentally conduce to—basic epistemic values. If there is such a thing as epistemic value that is derivative but not instrumentally so, there will be greater and lesser degrees of it. Suppose then that some action or attitude A instrumentally conduces most effectively to certain basic values, but does *not* derive any value in a *non-instrumental* way from more basic values. Plausibly, there will exist a counterpart of A—A*—which not only conduces instrumentally to these values, but *also* derives further value in a non-instrumental way from the more basic values. If so, A* is better than A. This is impossible on TELEOLOGY—VALUE ALIGNMENT. So, if there *is* non-instrumental value derivation, TELEOLOGY—VALUE ALIGNMENT is false. But (i) is just the contrapositive of this claim. So (ii) is true.

This was schematic. To flesh this out, let me get a variety of non-teleological value derivation on the table.¹⁰ Consider some facts. One can manage to promote a certain value even if one does not properly value this value. For example, consider a governor from one political party passing good legislation of which he does not personally approve in a state that supports causes that conflict with his beliefs. He does this to maintain an approval rating. In doing so, he promotes the right values, but his actions do not seem worthy because he does not himself *value* these values. A converse case is imaginable. A person's action can manifest proper valuing of a value even if she does not succeed in promoting it. For example, think of another politician who proposes some good legislation and fights hard to get it passed simply because it is correct, only to have an elected majority on the other side shoot down the cause. Even if he failed to promote the good, there was something good about the attempt. What? Plausibly, that it manifested a fitting attitude to value. He properly cared about the right stuff. This was what was lacking in the first example.

The best possibility would be a third: promoting genuine value in a way that manifests a fitting attitude to genuine value, and *because* of that fitting attitude. This affords a clear illustration of the kind of comparison between possible acts mentioned in my schematic argument. Our third possibility here is the A* that is better than A (the governor case).

3.1.2. *What Non-Teleological Value Derivation Could Be*

What the cases just considered illustrate is a *non-teleological* form of value derivation. It is not

¹⁰ See Sylvan (forthcoming) for further discussion of this.

plausible that having a proper attitude to value has *underived* value. Certainly, caring about value is itself valuable. But it is not as if there is no explanation of why it is valuable! It is valuable *because* caring is a fitting response to this value, and this value is, well, a real value. This fits with the fact that caring *itself*, considered independently of its object, seems to have *indeterminate* value. We need to know about the object valued to know whether the caring is good. The caring is good when it is *because* the object is good, and caring is a fitting response to that good.¹¹ Clearly, this derivation is not teleological. For as we saw in the second example above, simply having a fitting response to value does not ensure that this value will be promoted. It may not make it at all more objectively likely.

The example of value derivation that we're now considering is subsumed by this model:

Fitting Response Derivation: Necessarily, if R is the fitting response to a genuine value V (in some domain D), R derives value (in D) in virtue of (i) the fact that V is a genuine value (in D), and (ii) the fact that R is a fitting response (in D) to it.

Again, this is not a teleological form of value-derivation. This is clear given our example of the thwarted politician who sought to get the right legislation passed because he properly valued this legislation. This person's proper valuing did *not* result in the promotion of the value. Still, the caring was good, and good *because* it targeted the right thing. There is no instrumental explanation of why the caring was good. But there is the explanation given by Fitting Response Derivation.

One might, of course, try to add *caring about the good* as a basic value, and then claim that the politician did trivially promote this further basic value. But this is implausible, given the competing intuition that caring about the good is only good *because* it is directed at the good, and is a fitting response to that target. Caring about the good simply doesn't seem to be a *basic* value.

Unsurprisingly, I think this model of value derivation extends to epistemology (as I argue in Sylvan (forthcoming)). Indeed, I believe the model of Fitting Response Derivation is crucial to understanding how true belief could be the sole fundamental epistemic value. In brief, I think we must recognize (a) that having fitting attitudes to the epistemic value of truth (understood as the *standard* of belief formation) is epistemically good, but (b) that it is epistemically good *because* truth is itself epistemically good. (b) makes sense given Fitting Response Derivation. (a) is crucial to understanding why rationality and justification are both epistemically valuable, as we will see. For, as I will argue, responding to requirement of epistemic rationality and to epistemic reasons *is a way of fittingly responding to the epistemic value of truth (understood as the standard of belief)*.

This will all be unpacked at great length later; for now, I will just stress that the Fitting Response Derivation model is so popular outside epistemology that invocation of it in epistemology should be met with optimism.¹² But at present, we must see why we need such a model in epistemology at all. The model is non-teleological; the status of epistemic teleology is what is at issue.

11 This is not, of course, to assume that caring is a fitting response to *all* goods. The paradox of hedonism strongly suggests that it is more fitting *not* to care about pleasure, though it is good. Pleasure is a *merely* teleological good—one whose value consists in its being “to be promoted”.

12 See Hurka (2001) for discussions of its illustrious history of endorsement. Hurka extensively defends the account. He does claim that having a fitting response to the good is *intrinsically* good. But many have seen a crucial need to distinguish between basic and non-basic intrinsic value; cf. Feldman (2000) and Zimmerman (2001). Hurka would agree. He defends a *recursive* account of intrinsic value, on which there are certain base-level intrinsic goods (e.g., pleasure and knowledge), and then non-base-level goods generated by applying something structurally just like Fitting Response Derivation. What I think this shows is that value theory should be more interested in the distinction between basic and non-basic—or underived and derived—value than in the distinction between intrinsic and extrinsic value. If there can be exceedingly non-basic intrinsic goods, what should interest us more are the basic goods and the modes of value derivation generate non-basic ones.

3.1.3. Premises (ii) and (iii)

If epistemic teleologists could explain all the evaluative facts in epistemology without appealing to Fitting Response Derivation, we would not need this model. Unfortunately, however, epistemic teleologists *cannot* explain all the evaluative facts in epistemology. This is what premise (ii) tells us.

Let me defend it. The epistemic teleologist must explain the epistemic value of any non-basic epistemic value in instrumental terms. As I insisted and argued earlier, statuses like rationality and justification do not plausibly have basic epistemic value. So, given TELEOLOGY—VALUE ALIGNMENT, they must have instrumental epistemic value, relative to more fundamental epistemic values.

This, however, cannot be true. To build up to this, recall how we already saw in §2.5 that the only real option for the epistemic teleologist is to try to instrumentally explain the epistemic value of rationality and justification in terms of the deeper epistemic value of true belief. There simply isn't a plausible story on which the epistemic value of rationality (or justification, for similar reasons) is *instrumentally* derived from the epistemic value of knowledge or understanding. Of course, I also argued that there cannot be a clear *veritistic* explanation in all or even many possible cases of the epistemic value of rationality. But let's set that aside for the moment. Consider justification. If I was right to suggest in the last chapter that many epistemologists ought to distinguish conceptually between justification and rationality and view the former as having an *objective* connection to truth, the teleologist might think that some such story is to be had for justification. Perhaps that just shows that justification is what is epistemically valuable—rationality isn't.

Unfortunately, this cannot be right. The reason is a simple one familiar from the literature on the swamping problem. Recall Zagzebski's parable. The fact that a good cup of coffee was produced by the best coffee maker doesn't add anything to the value of this cup of coffee. This supports:

(Swamping) If the value of X is only instrumental relative to a further good and that good is already present in Y, it can confer no additional value to Y.

Suppose the value of justification or rationality really were instrumental relative to true belief. Then the Swamping Premise implies that a justified true belief or rational true belief cannot be better than a true belief. This, however, is clearly false. Which is to say that premises (ii) and (iii) are true.

3.1.4. What about Goldman (and Olson)?

Of course, some (e.g., Goldman and Olson (2009)) insist that there is a *broadly* instrumental explanation of why justified true belief is better than true belief. But these proposals involve such a switch in the meaning of “instrumental value” that I cannot see them as supporting TELEOLOGY—VALUE ALIGNMENT. Indeed, the plausibility of their core illustrations *refutes* TELEOLOGY—VALUE ALIGNMENT.

In explaining this proposal, Goldman (2012) relies on the following sort of analogy. We tend to value products produced by distinguished companies with a long history of success (e.g., Rolex) more than products produced by lesser known companies. In principle, however, *some* products by each company could be intrinsically the same. Goldman construes this as showing that an item can *instrumentally* derive value by being the output of a type of source that has great value in instrumental terms. Rolex is a better watch-making company, partly because of its reliability in producing fine watches. A given Rolex watch is then better because it is the product of such a company, *even if it is intrinsically the same* as an impressive watch that the lesser company produced by accident.

The derivation is what Goldman calls *type-instrumental* derivation. While a mode of value-derivation corresponding to the one observed by Goldman does exist, his terminology misleads. This is

not really a form of instrumental value derivation. Just reflect on the structure of the case. The *source* from which the watch derives its value is Rolex. Clearly, the watch is not valuable as a means to the valuable target from which it derives its value (Rolex). Nor is the *type* of which the watch is an instance instrumentally directed at the *type* of company that Rolex is (a reliable company). As far as instrumental relations are concerned, exactly the reverse is true! Rolex is a good company because it produces watches of this fine sort. While the value of the *source*—namely, the company Rolex—is instrumental, the *extra* value that the watch obtains from this source is not instrumental. That is exactly what the swamping problem shows! Sure, the watch does derive its value from something that has instrumental value—namely, the company Rolex. But we should not confuse this with the claim that the watch's value is *instrumentally* derived. In no familiar sense of 'instrumental' is this true. If anything, the way in which the watch derives its value from the instrumentally valuable source seems to be a *non-instrumental* form of value-derivation.

So, while Goldman's example is an illuminating one that may help epistemologists to understand how swamping is avoided, the example does not show that *teleologists* can avoid swamping. The example shows exactly the reverse. It shows that there is a form of value derivation that cannot be understood by the teleologist. For this sort of reason, I do not see Goldman and Olson's proposal to be of honest use to the defender of TELEOLOGY—VALUE ALIGNMENT.

Does this undermine Goldman's solution to the swamping problem? Not exactly, nor is it my intention to show that. I have argued that *teleologists* cannot appeal to this solution, and that Goldman's terminology is misleading. But reliabilists can *in principle* appeal to non-instrumental modes of value derivation. Remember what I marked at the outset in §1.1. While reliabilists like Goldman sometimes appeal to epistemic teleology to motivate their reliabilism, reliabilism itself is a first-order general criterial theory. As such, it is itself *compatible* with any number of foundational theories. Yet I am discussing *foundational* theories. What I've shown is that Goldman and other reliabilists should not accept a certain foundational theory—epistemic teleology. Unless it is *impossible* for a reliabilist to coherently reject epistemic teleology, this does *not* undermine Goldman's core commitments. It also does not show that reliabilists can't solve the swamping problem.

Perhaps they can solve it if they base their reliabilism on *non-teleological* foundations. This idea may seem wild. But it isn't. Consider how Parfit (2011) gave a *Kantian* argument for *rule consequentialism*. Kantianism is best understood as a *foundational* theory. Rule consequentialism is more like a first-order criterial theory. Whatever dialectical complaints one might lodge against Parfit, he is importantly right that Kantianism and rule consequentialism are *compatible*, even if they don't *coincide*. The point now is parallel. Process reliabilism is like rule consequentialism (cf. Goldman (1986)). Just as rule consequentialism is compatible with Kantian foundations, so process reliabilism is compatible with non-teleological (perhaps Kantian!) foundations.

3.2. *The Rest of the Scanlonian Argument*

I have now defended the key premises in my argument against TELEOLOGY—VALUE ALIGNMENT. Given this argument, the first premise of the Scanlonian Argument holds. But the rest of the Scanlonian Argument must be defended. Fortunately, this task is not difficult, given my discussion about the relationships between the different varieties of epistemic teleology in §1.3.

Having defended the first premise, there are two premises that need defense (the others follow):

2. If TELEOLOGY—VALUE ALIGNMENT is false, then TELEOLOGY—DOXASTIC RIGHTNESS is also false.
4. If TELEOLOGY—DOXASTIC RIGHTNESS is false, then TELEOLOGY—GOOD REASONS is false.

I will take these in reverse.

The case for (4) is straightforward. Remember that TELEOLOGY—GOOD REASONS can be viewed as a consequence of two theses:

(RIGHTNESS CONSTRAINT—REASONS) R is a good epistemic reason to form a doxastic attitude D with respect to a proposition P at *t* if and only if R's truth at *t* would make it objectively more likely that D-ing with respect to P would be epistemically right D-ing than that it would be epistemically wrong D-ing.

(TELEOLOGY—DOXASTIC RIGHTNESS) Forming a doxastic attitude D with respect to P epistemically rightly is a matter of in fact promoting basic epistemic values in virtue of D(P)-ing (and not basic disvalues).

RIGHTNESS CONSTRAINT—REASONS should be uncontroversial. If this assumption is true and (TELEOLOGY—DOXASTIC RIGHTNESS) is false, then (TELEOLOGY—GOOD REASONS) is false too. After all, if (TELEOLOGY—DOXASTIC RIGHTNESS) is false, then there will be conditions on right D(P)-ing other than goal promotion. If so, imagine a case in which the fact that R massively increases the objective probability of goal promotion in virtue of D(P)-ing, but massively decreases the objective probability that other conditions on rightness are satisfied—indeed, imagine it *guarantees* they will not be satisfied. Then it follows by RIGHTNESS CONSTRAINT—REASONS that R cannot be a good reason for D(P)-ing. Since it *independently* follows from (TELEOLOGY—GOOD REASONS) that R is a good reason, the falsity of (TELEOLOGY—GOOD REASONS) follows from the falsity of (TELEOLOGY—DOXASTIC RIGHTNESS), given RIGHTNESS CONSTRAINT—REASONS. Since RIGHTNESS CONSTRAINT—REASONS is true, premise (2) in the Scanlonian Argument is true.

This was an abstract argument, so consider an analogy with ethics to understand it concretely. Suppose there are conditions on rightness besides *promotion* of the good—*restrictions*, let's say. Some fact might make it objectively certain that some act would promote the good, but also guarantee that a restriction is violated. If so, that fact couldn't be an ethically good reason for doing that act, plausibly. An ethically good reason has to increase the objective probability that some act would be right. But the fact at issue here guarantees that the act isn't right. What we've just shown is precisely that denying teleology about the rightness of acts leads, given a natural connection between rightness and reasons, to denying teleology about reasons. Since this is parallel to my argument, it is a useful way of understanding it. Of course, it is an open question *what* the analogue of a restriction is in epistemology. We will get to that—that is a substantive matter. The point now is just that *if* there are restrictions, (TELEOLOGY—GOOD REASONS) is false. So, again, (2) in the Scanlonian argument holds.

What about (4)? The argument for (4) is exactly structurally like the argument for (2). Remember that TELEOLOGY—DOXASTIC RIGHTNESS can be viewed as a consequence of two theses:

(EPISTEMIC VALUE CONSTRAINT—EPISTEMIC RIGHTNESS) D(P)-ing epistemically rightly is a matter of forming the epistemically best attitude available with respect to P.

(TELEOLOGY—EPISTEMIC VALUE ALIGNMENT) Forming the best doxastic attitude D with respect to P is a matter of in fact most efficiently promoting basic epistemic values in virtue of D(P)-ing.

Once we realize that value itself needn't be understood teleologically, the first thesis loses its air of controversy. Obviously, it would be controversial to assume that the right covaries with the good in the sense *consequentialists* have in mind. But what they have in mind is stronger: the right covaries with facts about the *promotion* of the good. Deontologists can deny this without denying the generalized version of the first thesis. They can agree that the right act is the best act. But they can insist that this claim doesn't support consequentialism, by insisting that the best act is not necessarily the act that promotes value most efficiently.

Once this distinction is borne in mind, I think we ought to accept the first thesis for what it is—namely, a platitude. But if it is true and TELEOLOGY—EPISTEMIC VALUE ALIGNMENT is false, then it is easy to argue that TELEOLOGY—DOXASTIC RIGHTNESS is false. After all, if TELEOLOGY—EPISTEMIC VALUE ALIGNMENT is false, then there will be conditions on what makes for *best* D(P)-ing other than overall optimal goal promotion in virtue of D(P)-ing. If so, we can simply imagine a case in which D(P)-ing would be perfectly efficient with respect to goal promotion in virtue of D(P)-ing but *terrible* with respect to the other parameters. Then it will follow by EPISTEMIC VALUE CONSTRAINT—EPISTEMIC RIGHTNESS that D(P)-ing here would not be right D(P)-ing. Since it *independently* follows from TELEOLOGY—DOXASTIC RIGHTNESS that D(P)-ing here *would* be right D(P)-ing, TELEOLOGY—DOXASTIC RIGHTNESS, the falsity of TELEOLOGY—DOXASTIC RIGHTNESS follows from the falsity of TELEOLOGY—EPISTEMIC VALUE ALIGNMENT, given EPISTEMIC VALUE CONSTRAINT—EPISTEMIC RIGHTNESS. Since EPISTEMIC VALUE CONSTRAINT—EPISTEMIC RIGHTNESS is true, premise (4) in the Scanlonian argument is true.

3.3. *Taking Stock*

I cannot see clear objections to my defenses of premises (2) and (4) in the Scanlonian Argument. Since I've also rebutted the salient objections to my defense of premise (1), this argument would appear to succeed. So, we ought to reject TELEOLOGY—GOOD REASONS, and precisely because TELEOLOGY—EPISTEMIC VALUE ALIGNMENT fails. TELEOLOGY—EPISTEMIC VALUE ALIGNMENT fails because it leads inevitably to swamping problems about the value of justification (and rationality—though its value is currently under dispute). A justified true belief really is better than a true belief, but TELEOLOGY—EPISTEMIC VALUE ALIGNMENT makes this impossible to explain. While some teleologists *claim* that it can be explained (e.g., Goldman), I have argued that they do not succeed in showing that it can be explained on grounds that are *genuinely teleological*. Indeed, the grounds to which they point undermine TELEOLOGY—EPISTEMIC VALUE ALIGNMENT. While this may not—as I've noted—be problems for them *qua* reliabilists, that is simply because reliabilism *per se* doesn't imply teleology. Like rule consequentialism, it could in principle be given non-teleological foundations.

First, however, I am going to turn to explain in §4 how the same variety of non-teleology about epistemic value discussed earlier in §3.1.2 both (i) establishes that the value of justification and of rationality is not swamped by the value of true belief, and (ii) gives us the basis for explaining why epistemic rationality has robust normative significance. The strength of our vindication of the normativity of epistemic rationality will depend on the precise version of non-teleology we accept. I will discuss the varieties of non-teleology in §5, turn to some conclusions and implications in §6.

4. **Vindicating the Value and Normative Significance of Rationality Non-Teleologically**

Remember that the reason why TELEOLOGY—VALUE ALIGNMENT fails is that there *must* be a non-teleological form of value derivation. If there isn't, then, as I argued, there will inevitably be swamping problems about the epistemic value of *justification*, not just of rationality. What I want to argue now is that the same type of story that helps us explain, non-teleologically, why justification is epistemically valuable *also* helps us to explain both (i) why rationality is epistemically valuable, and (ii) why rationality necessarily has at least *pro tanto* robust normative significance. I'll take these tasks in turn.

4.1. *Fitting Response Non-Teleology and the Unswamped Value of Epistemic Justification*

It will be helpful to consider first how the model works for justification, and obviates swamping. Because my task earlier was just to convince the reader that a non-teleological explanation is needed, I

did not lay out this explanation. That would have been premature. Let me now lay it out.

Earlier we considered examples from the practical case that made this model plausible:

Fitting Response Derivation: Necessarily, if R is the fitting response to a genuine value V (in some domain D), R derives value (in D) in virtue of (i) the fact that V is a genuine value (in D), and (ii) the fact that R is a fitting response (in D) to it.

One need only consult Hurka (2001) to see that this model has been widely endorsed throughout the history of ethics. Of course, the model that I need is a special case, since I am doing epistemic value theory, not value theory *simpliciter*:

Fitting Response Derivation—Epistemic: Necessarily, if R is the fitting response to a genuine epistemic value V, R derives epistemic value in virtue of (i) the fact that V is a genuine epistemic value, and (ii) the fact that R is an epistemically fitting response to it.

To get this model to apply, we need a plausible candidate for a fundamental epistemic value. On my view, truth (in belief) is the crucial candidate. Because I am not a teleologist, I do *not* think that we should understand the epistemic value of truth in belief to consist in the purported fact that we ought epistemically to promote the *state* of true belief. Rather, I propose that we should understand the epistemic value of truth in belief to fundamentally consist in our having an obligation to fittingly respond to the *standard* of truth in belief, which includes *rules* like: believe P only if P is true.

We can, I believe, explain the value of the *state* of true belief in a straightforward way on this model. One of the ways to fittingly respond to the standard *believe P only if P* is to *comply* with this standard. But one of the ways of complying with this standard is to have a true belief. Accordingly, having a true belief will be epistemically valuable, given *Fitting Response Derivation—Epistemic*. So, although I do not believe that the *fundamental* epistemic value is the state of believing truly, I do think this the *pro tanto* epistemic value of believing truly is easily derived from the epistemic value of the *standard* of truth in belief. Notice: what it *is* for the standard to be valuable is for it to be something that we ought to respond to in certain ways. While some talk of “values” simply is talk of concrete objects, I believe that there is a familiar sense in which something abstract like a *standard* can be a *value*: it can be something that is *properly valued*. To be a value *just is* to be a proper object of valuing. Since standards can be proper objects of some varieties of valuing, standards can be properly valued.

We can also explain the value of justification *via Fitting Response Derivation—Epistemic*:

1. Fittingly valuing true belief requires responding to sufficient epistemic reasons. (Assumption.)
2. If (1) is true and true belief is a fundamental epistemic value, responding to sufficient epistemic reasons is necessarily derivatively epistemically valuable *pro tanto*. (*Fitting Response Derivation—Epistemic*.)
3. True belief is a fundamental epistemic value. (Assumption.)
4. So, responding to sufficient epistemic reasons is necessarily derivatively epistemically valuable *pro tanto*. (1-3)
5. But doxastic justification consists in responding to sufficient epistemic reasons. (From Ch. 1)
6. So, doxastic justification is necessarily derivatively epistemically valuable *pro tanto*. (4-5)

This is a non-teleological explanation of the epistemic value of doxastic justification. How does it help

us to avoid swamping? Well, Fitting Response Derivation contravenes the following thesis:

Generalized Swamping. If the value of a property possessed by an item is only *non-fundamental* value relative to a further good and that good is already present in that item, it can confer no additional value.

This is “*Generalized*” Swamping because it replaces 'instrumental' with 'non-fundamental'. And it is false. For it is better if a situation has the property of being one in which a fitting favorable response to X's value exists than it would be if it *only* had the property of being one in which X exists. Since the former property has only non-fundamental value relative to the latter, Generalized Swamping fails.

An example in which one has a justified true belief is an example of this general phenomenon. The epistemic value of justified belief is derived from the epistemic value of true belief *via* Fitting Response Derivation, *via* the argument just offered. But it's a general fact that when a value derives value from X in this non-teleological manner, its value is not swamped by the presence of X. This is why the epistemic value of justification is not swamped by the epistemic value of true belief.

What about the argument itself? (1) seems plausible on its face. If you fail to respond to sufficient *epistemic* reasons—which are, after all, considerations that count for or against the *truth* of beliefs—it certainly seems that you do not properly value true belief. (2) is just a consequence of Fitting Response Derivation. (3) is, I think, defensible—and widely accepted. (4) is a consequence. (5) was established in Chapter 1. So, this argument would seem to stand—though (1) will be discussed further once we've seen the argument for the unswamped epistemic value of rationality.

4.2. *Fitting Response Non-Teleology and the Epistemic Value of Rationality*

Exactly the same kind of argument can be used to show that epistemic rationality is epistemically valuable. Let's build up to it informally. Note that it is plausible that if you fail to respond to what *clearly appear* to be epistemic reasons, you simply cannot fully properly value the standard of truth in belief. There is something manifestly lacking in one's valuing of truth if one is so constituted that *clear apparent epistemic reasons* against one's beliefs would “leave one cold”. This seems *even more* plausible than the parallel claim about genuine as opposed to *apparent* epistemic reasons. Perhaps one might be excused if one fails to see *that* certain reasons are relevantly genuine epistemic reasons, even if they really *are* genuine epistemic reasons. If so, then while one may not in some sense bear the right relation to truth, it is not as if one would fail to fully properly value true belief. A clear charge against one's ability to value truth would be warranted if one failed to respond to clearly apparent reasons.

This leads to an even clearer argument for the epistemic value of rationality from *Fitting Response Derivation—Epistemic*:

7. Fittingly valuing true belief requires responding to apparent sufficient epistemic reasons. (Assumption.)
8. If (1) is true and true belief is a fundamental epistemic value, responding to apparent sufficient epistemic reasons is necessarily derivatively epistemically valuable *pro tanto*. (From *Fitting Response Derivation—Epistemic*.)
9. True belief is a fundamental epistemic value. (Assumption.)
10. So, responding to apparent sufficient epistemic reasons is necessarily derivatively epistemically valuable *pro tanto*.
11. But doxastic rationality consists in responding to apparent sufficient epistemic reasons. (From Ch. 1)

12. So, doxastic rationality is necessarily derivatively epistemically valuable *pro tanto*.

This argument is even more decisive than the last, since (7) is even more plausible than (1) in the other argument. This is for the reason recently mentioned. One might be able to imagine someone properly valuing truth as a standard but being “left cold” by *objective* epistemic reasons, simply because this person is *in the dark* about these reasons. But it is extremely hard to make sense of someone who values truth as a standard while being “left cold” by *apparent* epistemic reasons—that is, considerations which *appear to be* epistemic reasons from this person's perspective (in either a doxastic or non-doxastic sense). So (7) is hard to deny. The rest of the moves in this argument are defensible for the same reasons as the moves in the earlier argument. Indeed, they are essentially the same moves.

While the epistemic value of doxastic rationality is derived from the epistemic value of truth, on this account, swamping is nevertheless avoided. This is again because, unlike instrumentally derivative value, value that is derived *via* Fitting Response Derivation cannot be swamped by the presence of the deriving value. It is obviously (epistemically) better to have a fitting response to (epistemic) value *and* for that value to be promoted than *merely* for that value to be promoted. Since the value of a fitting response to value is derived from that value, given Fitting Response Derivation, this is simply a counterexample to Generalized Swamping. Swamping is local to instrumentally derived value.

4.3. *Can the Robust Normative Significance of Rationality Be Vindicated?*

There are a couple of ways of seeing why the ideas connected with Fitting Response Derivation could help to vindicate the robust normative significance of rationality, and not just its value. The first way leads only to a partial vindication. The second way deliberately has a false premise, since the reasons why the premise is false will bring out the further commitments that would be needed to *fully* vindicate the robust normative significance of rationality.

The simplest but also weakest way is to argue from (VALUE ALIGNMENT CONSTRAINT—RIGHTNESS) and (RIGHTNESS CONSTRAINT—REASONS), granting ourselves the conclusions of the last section. Specifically, we can argue as follows:

- A. Doxastic rationality is derivatively epistemically valuable, and in a way that isn't swamped by the value of true belief.
- B. So, for a belief to be epistemically *best*, it must be doxastically rational.
- C. (VALUE ALIGNMENT CONSTRAINT—RIGHTNESS) implies that one believes P rightly only if believing P would be epistemically best.
- D. (RIGHTNESS CONSTRAINT—REASONS) implies that R is a good epistemic reason for believing P only if R's truth would make it more likely that believing P would be epistemically right.
- E. So, given that these constraints are both correct, R is a good reason for believing P only if R's truth would make it more likely that believing P would be epistemically best.
- F. But it follows from (E) and (B) that for R to be a good reason for believing P, R must at least make believing P rational, since that's required for believing P to be epistemically best.
- G. Therefore, an epistemic reason cannot be good unless it is also makes a belief rational.

(G) is weaker than what we want in two respects. We want considerations that make a belief rational to always be good *pro tanto* reasons. Indeed, we want sufficient rationalizers to be sufficient reasons. (G) establishes only that a reason is good *only if* it is a rationalizer. It does not establish that a reason is good *if* it is a rationalizer. Plausibly, however, we can still get the first stronger conclusion from (G), and *partially* vindicate the robust normative significance of rationality. For it is plausible that a reason is a *pro tanto* good reason to believe P if this reason guarantees that believing P would always be *better* other things being equal. But rationalizers *do* have that property, given the argument from the last subsection. So rationalizers are always good *pro tanto* reasons for belief.

Crucially, this much could not be obtained from any plausible version of epistemic teleology, so the *modest* version of fitting response non-teleology we are now considering makes a big advance. As we will see, a *full* vindication of the normativity of rationality will require an explicitly stronger view.

Let us consider a different argument for the stronger conclusion. This argument is not meant to be decisive. It is meant to reveal the further elements that a version of non-teleology would need to have if it is to *fully* vindicate the robust normative significance of epistemic rationality. Here it is:

- i. If R is a fitting response to V and V is a fundamental epistemic value, there is always sufficient epistemic reason to have R to V.
- ii. Fittingly valuing true belief requires responding to apparent sufficient epistemic reasons.
- iii. True belief is a fundamental epistemic value.
- iv. So, there is always sufficient epistemic reason to respond to apparent sufficient epistemic reasons.
- v. Doxastic rationality consists in responding to apparent epistemic reasons.
- vi. So, there is always sufficient epistemic reason to be doxastically rational.

Premises (ii), (iii) and (v) overlap with our earlier argument in §4.2. So if that argument is sound—as it seemed to be—we should accept these premises too. (iv) and (vi) are consequences. So the only premise one could question is (i). There are some reasons (i) fails, and hence why it and the other premises need to be strengthened. This will lead naturally to the next section, in which I will catalog some different versions of Fitting Response Non-Teleology. One of these versions, if true, would yield a sound argument of this form with the same conclusion.

Here are the two reasons why (i) cannot be true. Firstly, there might be a fitting response to V *besides* R, and this fitting response might in some cases be *incompatible* with R-ing. If so, it is not obvious that there would be a *sufficient* reason to R, since there might be just as good reason to have the other response to V, or better reason. Secondly, there might be some *different* value V* to which there is a different response that, in a given case, is *also* incompatible with R-ing. If V* is more important than V, then although R-ing would still be fitting to V, there would not be sufficient reason to R(V). For R(V)-ing is incompatible with the fitting response to V* and V* is more important.

To fully vindicate the robust normativity of epistemic rationality in the way at issue here, both of these possibilities will have to be ruled out. Moreover, *if* we rule out these possibilities, we *can* fully vindicate the robust normativity of epistemic rationality. For the following revised argument seems sound if the revised versions of (i) and (iii) are true:

- i*. If R is the most (or only) fitting response to V and V is the sole fundamental (or most important)

epistemic value, there is there is always sufficient reason to have R to V.

ii*. The most (or only) fitting response to the epistemic value of true belief is to respond to apparent sufficient epistemic reasons.

iii*. True belief is the sole fundamental (or most important) epistemic value.

iv*. So, there is always sufficient epistemic reason to respond to apparent sufficient epistemic reasons.

v*. Doxastic rationality consists in responding to apparent epistemic reasons.

vi*. So, there is always sufficient epistemic reason to be doxastically rational.

Ultimately, I think that this is the right argument. In the next chapter, I will defend (i*) and (iii*). For the moment, however, I am noting that this is a version of the most obvious way to fix the earlier argument to get a full vindication of the robust normativity of epistemic rationality. I say “a version” because I am uncertain that there is no plausible alternative fundamental epistemic value that could be substituted for “true belief” in a sound variant of this argument. But to keep things simple and candid, I am discussing the argument that reflects the view I will defend at the end of the day.

5. Varieties of Non-Teleology and Strengths of Vindication

We saw abstractly in §4.3 that there have to be several varieties of fitting response non-teleology, individuated by (i) how they understand what it is to fittingly value epistemic value, and (ii) their assumptions about what the basic epistemic value(s) might be. For simplicity, I will examine veritistic views, and focus on element (i) and further elements that underlie the choice constituted by (i). I will return to variations on (ii) in the next chapter. The discussion §4 was deliberately abstract, since I wanted us to see clearly what *form* these views would have to have, to get vindications of the normativity of rationality of different strengths. It is now time to put some flesh on this discussion.

5.1. *Mere Non-Teleology*

What makes a view about epistemic value teleological is its viewing epistemic values as being epistemic goals, period. In Scanlon's words, these are views on which to be epistemically valuable just is “to be promoted” as an epistemic end. The negation of this view doesn't entail that epistemic values are not goals *at all*. It only need involve the claim that there is a proper way of valuing these values that is not teleological. Some epistemic values are not *merely* epistemic goals. This is compatible with its being *prima facie* epistemically proper to value them as goals too.

The comparison with ethics makes this obvious. Those who reject teleology in ethics do not claim that promoting value doesn't matter at all, in any case. That would be crazy. There are some practical values—e.g., pleasure—that are best understood teleologically. We have no obligation to care explicitly about pleasure. Given the paradox of hedonism, we probably shouldn't, but just get more instances of it into our lives indirectly. Those who reject teleology in ethics just insist that there are some values whose value does not consist in the fact that we ought to cause more instances of them to exist. While it may be plausible that pleasure is something to be promoted, it is implausible that friendship's value consists in its being “to be promoted”. It is plausible that properly valuing friendship does not consist simply in trying to have as many friendships as possible. It also requires valuing one's

friends as friends by participating in the activities constitutive of being a good friend. That isn't to say that having more friendships isn't also a good thing, but the fittingness of such promotion does not *exhaust* the value of friendship. Similarly, it is implausible that properly valuing art consists just in seeing to it that more art exists, and that existing art persists. Someone who is left cold by art and who is incapable of aesthetic appreciation could do this, while failing fully to value art.

One kind of epistemic non-teleology might view the epistemic value of true belief in this way. It might concede that promotion is *an* epistemically fitting response to the epistemic value of true belief but deny that true belief's being epistemically valuable is *exhausted* by this fact. We also ought to properly *value* the standard of truth in belief by complying with apparent epistemic reasons. Failing to properly value truth in this way would be to miss out on a very important part of what it is for true belief to be epistemically valuable. Part of its epistemic value is as a standard in doxastic deliberation.

If one accepts this kind of view, it is unobvious that a *complete* vindication of the normativity of epistemic rationality is possible. After all, promotion and non-teleological responses like valuing-as-standard can conflict. Clearly, as we saw from the start, there are worlds where complying with apparent epistemic reasons would systematically ensure that one would have more false beliefs than true beliefs. If true belief *also* ought to be promoted, it is unclear that rationality *would* have decisive robust significance on a view of this form. By analogy, consider art again. Suppose that by acting in some way W, I could cause an enormous amount of great art to exist, and also cause all existing art to be preserved until the end of time. But by acting in way W, I would also ensure that only a few people could ever appreciate this art, and deprive most people of any ability to appreciate art. On the other hand, I could act in way W* and cause all the art in the world except that in the MoMA to cease to exist. Yet I could also enhance everybody's capacity to appreciate this art maximally. I could ensure that everybody gets a lifelong membership to the MoMA, a good education in modern art, and the resources to visit Manhattan often. Suppose finally that W and W* are the only options. It is unclear that W* is the right one. Maybe promotion can trump appreciation as a fitting response to art's value.

So if we are searching for a *complete* vindication of the normativity of rationality, we must defend some further claims. We need a defense of the claim that promotion is never the most epistemically fitting response to the epistemic value of true belief. If one is willing to honor teleological intuitions at all, I am uncertain that one will be able to have a stable defense of this position. I will return to this in the next chapter.

5.2. *Anti-Teleology without Full Vindication (and with Full Vindication)*

One could adopt a more radical view, and deny that promotion is necessarily a fitting response to the epistemic value of true belief. Of course, it will not be plausible on any view that it is *never* a fitting response. Suppose we know that P entails Q, know that P, and are considering the question whether Q. Here we are rationally permitted to believe a truth (namely, that Q), and hence permitted to promote the value of true belief by causing a true belief to exist. What is permissible is fitting. So it must *sometimes* be fitting to promote true belief. What a plausible form of anti-teleology ought to say is that promotion is not a *fundamentally* fitting response to the epistemic value of true belief. In this case, for instance, the anti-teleologist could say that it is *because* we are rationally permitted to believe the true proposition Q that it is fitting to promote true belief. A plausible form of anti-teleology ought then to add that promoting true belief (by causing more true beliefs to exist) will *only* be fitting if it is *made* fitting by deeper, non-teleological ways of valuing—e.g., the way of valuing that requires complying with apparent reasons. Since in some cases we are rationally required to believe false propositions, these are cases in which promotion would *not* be a fitting way of valuing true belief.

Does anti-teleology so understood give us a full vindication of the robust normative significance

of rationality? Not by itself. For one could hold the following conjunction of views:

One fittingly values truth if and only if one complies with apparent epistemic reasons that are also genuinely good epistemic reasons.

It is possible for apparently good epistemic reasons to be *merely* apparently good.

Complying with truly good epistemic reasons needn't ensure true belief justification is not plausibly factive (*pace* Littlejohn (2012)). So this view could be a version of anti-teleology, if it claimed that what it is for truth to be valuable is to be a fitting object of valuing in the sense spelled out by the first claim. Nevertheless, this is *not* a view on which the epistemically best thing to do is to comply with requirements of rationality. For one is rationally required to respond merely to *apparent* reasons; one isn't irrational if one doesn't align one's attitudes with objective facts about how good these are.

This view is simply listed as a *dialectical possibility*. Whether this view is *stable* is an open question. In the next chapter I'll argue that this is unstable. For now, however, the intelligibility of this view suggests that being an anti-teleologist isn't sufficient for a full vindication of the normativity of rationality. So there is more work that one must do to establish a full vindication.

This also shows us what a version of anti-teleology that enables full vindication would have to look like. This view would have to reject the second claim, or perhaps claim that one fittingly values truth iff one complies with apparent epistemic reasons, full stop.

5.3. *Aren't There Other Varieties of Non-Teleology besides Fitting Response Non-Teleology?*

My presentation so far may convey the impression that the only way to be a non-teleologist who can vindicate the normativity of rationality is to exploit the idea that there are fitting ways of valuing epistemic value besides promotion. This may seem arbitrary. But while there are conceivable further views, I think the only *plausible* alternatives are simply instances of this approach. Why?

Let me start by ruling out a class of alternative I regard as implausible. To see why I regard this family of views as implausible, I will stress again that there is a strong conceptual connection between value and rightness, captured by VALUE ALIGNMENT CONSTRAINT—RIGHTNESS. Given this connection—which strikes me as nonnegotiable—I cannot make sense of views on which A-ing can be right in the objective sense used throughout while A-ing isn't best (or doesn't meet some satisficing threshold). As I have insisted throughout, the way to reject teleology is *not* to reject this conceptual connection between value and rightness. It is to reject the teleologist's implicit picture of value.

There is a tradition on which deontological views *just are* views on which facts about value fail to fix facts about rightness and what reasons there are. These people assume that the distinctive idea behind teleology is the idea that the good is prior to the right. This is confused. One can coherently think that value is the most basic normative concept and deny that to be valuable is to be “to be promoted”. On a view of this kind, it can be right to do what would not best *promote* value. But such a view does not sever the connection between value and rightness. So merely claiming that the good is prior to the right doesn't establish consequentialism or anything like it. Only the view that the goodness of states of affairs is prior to the right could establish this conclusion, when good states of affairs are understood as things that are “to be promoted” as goals.

Because I am unhappy with traditional ways of carving up the difference between teleological and non-teleological views, I do not see deontology traditionally conceived as an option. It involves the radical claim that A-ing can be right while A-ing isn't best (or doesn't meet some satisficing threshold). I find this at best confused, and at worst unintelligible.

Of course, traditional deontology is hardly the only *apparent* non-teleological alternative to the one I've promoted. Is there not also virtue theory? What about Kantianism?

I see these theories as subsumed by the style of non-teleology I've promoted. This is because of a further conceptual connection I regard as nonnegotiable. I think that values are necessarily items that are fittingly valued, and that whatever is fittingly valued is a value. Given that I also think that there is a necessary connection between value and rightness (and reasons too), I think that good positions on what epistemic rightness involves and what makes epistemic reasons good will require a view to the effect that there are epistemically fitting non-teleological ways of valuing epistemic value.

This comports well enough with those who might adopt a Kantian conception of the fundamental epistemic norm—say, as the norm to *respect* the standard of truth in belief. Can it also comport with virtue epistemology? I think so. We can simply take the manifestation of epistemic virtue to be a way of fittingly valuing some choice epistemic value. On conceptions of virtue outside epistemology, this is, indeed, a popular idea.¹³

Might the virtue epistemologist claim that this involves overintellectualization? Not plausibly, if she isn't implicitly overintellectualizing the business of fittingly responding to value. I understand this notion broadly, so that manifesting innate virtues or virtues acquired through early childhood development can count as fittingly valuing truth. It is not as if children who respond to perceptual evidence, testimonial evidence, and so on, *fail* to value truth as epistemic agents. On plausible accounts outside of epistemology, attitudes of proper valuing are often seen as dispositions to have certain kinds of positive responses to perceptions of, among other things, the presence or absence of the goods that they take as their objects. There is certainly no reason to doubt that people at younger stages of epistemic growth can possess attitudes of proper valuing in this sense.

Accordingly, such examples at best show only that the idea of valuing truth should not be overintellectualized. When it is made clear that this can be understood weakly, it seems best on the whole to embrace the connection between value and fitting ways of valuing and view non-teleological virtue theories as either *being* versions of fitting response non-teleology or as requiring implicit commitment to this view. At any rate, if they are to help vindicate the normativity of rationality, they *must* involve such a commitment given the connections I've insisted are nonnegotiable—viz., the connection between value and fitting ways of valuing, VALUE ALIGNMENT CONSTRAINT—RIGHTNESS, and RIGHTNESS CONSTRAINT—REASONS. These connections mandate fitting response non-teleology, even if it is concealed by talk of virtue, responsibility, or the like.

6. Concluding Remarks

It is time for a recap and some concluding remarks.

This chapter has accomplished three big things. First, we saw in §2 that epistemic teleology about norms encourages a skeptical view about the normative significance of epistemic rationality. This skepticism cannot be plausibly avoided by subjective teleology, unless one finds promising a solution as apparently trivial as: “Epistemic rationality is normatively significant because we ought rationally to be rational.” Subjective teleology cannot offer more than this or some merely verbal variation on it. This skepticism also cannot plausibly be avoided by exchanging truth for some other plausible fundamental epistemic value. It could be trivially avoided if epistemic rationality were a fundamental epistemic value, but this is implausible. Although epistemic rationality may be *non-instrumentally* epistemically valuable, it must be seen as *derivatively* valuable on any plausible view.

That brings us to the second thing this chapter has accomplished. We saw that teleology about

¹³ See Hurka (2001: 23-28) for a discussion of historical antecedents and (2001: 20) for a definition of virtue on which virtue simply consists in having fitting pro-attitudes to what is intrinsically valuable.

epistemic *norms* must be false if teleology about epistemic *value* is false, and that teleology about epistemic value *is* false. Teleology about epistemic value is false because it simply cannot be true that all non-fundamental epistemic value is instrumental epistemic value, in any natural sense of “instrumental”. This is because teleology about non-fundamental value leads inevitably to swamping problems about the value of justification (and rationality). Once we reject teleology about epistemic value because it leads to these problems, we can argue against teleology about epistemic norms by accepting the same assumptions that the teleologist about epistemic norms must herself accept—viz., (RIGHTNESS CONSTRAINT—REASONS) and (VALUE ALIGNMENT CONSTRAINT—RIGHTNESS). Since these assumptions are indeed plausible, teleology about epistemic norms fails.

The last thing that this chapter accomplished was twofold. Firstly, we saw that the model of epistemic value derivation that most clearly resolves the swamping problems about the epistemic value of justification also leads to a vindication of the epistemic value of rationality. Indeed, the vindication of the unswamped epistemic value of rationality is *even clearer* than the vindication of the unswamped epistemic value of justification. We also saw that this fact leads, given (RIGHTNESS CONSTRAINT—REASONS) and (VALUE ALIGNMENT CONSTRAINT—RIGHTNESS), to an at least partial vindication of the robust normative significance of epistemic rationality. But we also saw, secondly, that whether a *full* vindication is available will depend on the particular version of fitting response non-teleology one accepts. It is only obvious that an *anti*-teleological view would yield a full vindication.

It is worth concluding with implications for some mainstream epistemologists. Some reliabilists have indeed been veritist epistemic teleologists. They should not accept this view if they want epistemic rationality to matter as such (*pro tanto* or otherwise). Given, as we saw in Chapter 4, that epistemic rationality *must* matter as such at least *pro tanto* to vindicate the reliabilist's views about *defeat*, she must reject epistemic teleology. While it is unclear that this compromise is stable, this isn't an incoherent view. As Parfit showed in the ethical case, rule consequentialism could be coherently given Kantian rather than properly teleological foundations. Since process reliabilism *a la* Goldman (1986) is closely analogous to rule consequentialism, one should expect the same kind of coherent combination to be possible in epistemology. Whether this coherent combination is a *plausible* combination on the whole is another matter, to be discussed in the next chapter.

Those are some implications for one brand of mainstream externalism. The implications for internalism are also important. As I argued in the last chapter, even the most minimal form of internalism—namely, Supervenience Internalism—mandates collapsing the distinction between apparent and genuine epistemic reasons, and hence the distinction between rationality and justification. Given, as we saw, that this could only be true if one accepts fitting response *anti-teleology*, and not merely fitting response *non-teleology*, all internalists have some strong further commitments to defend.