

What the objective standard is good for

Julia Driver

Forthcoming in *Oxford Studies in Normative Ethics* (2), ed. Mark Timmons

A recent spate of articles has reinvigorated discussion of objective versus subjective standards of ‘rightness’ and ‘rationality.’ In a recent work I defended an objective standard of ‘right’ in consequentialist moral theory, by holding that the objective standard is prior to the subjective standard. However, the objective standard has been criticized as setting an unrealistic standard, one that, in effect, conflates rightness with good moral luck. What good is a standard of right that we cannot – and, in some cases, should not try to – live up to? This worry, on my view, is traceable to the assumption that the work ‘right’ does – indeed, the only work it does – is as a mode of evaluation in holding people responsible for what they do. However, I believe a case can be made for the practical significance of the objective standard in that *acceptance of* such a standard offers a way of measuring success. The objective standard allows for a very natural taxonomy of how our actions go morally wrong, a taxonomy that does, indirectly, provide guidance.

The account:

What counts as an *objective* standard? Typically, objectivity is understood as holding regardless of an agent’s psychological states – their beliefs, desires, and so forth. But this underdetermines the contrast between the objective and subjective consequentialists, as I have argued elsewhere.¹ Another characterization is that the standard, to be objective, is not in any way evidence-sensitive. This is distinct from the previous way of carving the distinction since evidence might bear on an agent’s beliefs, though the agent is not at all

¹ *Consequentialism* (Routledge, 2012).

aware of the evidence. For example, Peter Graham has recently defined an objective moral theory in contrast with an evidence-subjective theory:

A moral theory *T*, is *evidence-subjective=def.* according to *T*, necessarily, a person has the moral obligations that she has at a time solely in virtue of facts about her evidential situation at (or prior to) that time..

A moral theory *T*, is *objective= def.* it is not the case that *T* is evidence-sensitive.²

Graham is not discussing consequentialist theory *per se*, but the distinction would apply.

Often, evidence-subjectivity, or evidence-sensitivity, includes the agent's actual belief and/or desire states, either present beliefs or merely dispositional ones, but it need not. A person, for example, who thought that the right action was the action that the agent ought to believe maximizes the good, given the evidence available at the time, is proposing an evidence-sensitive standard that does not appeal to the agent's *actual* beliefs. Indeed, what makes the action right is not recoverable at all from the agent's own psychology at the time of action. For this reason, I have argued elsewhere that such 'subjective' theories do not solve a major problem that subjective theories are supposed to solve – the action ownership problem discussed by Frank Jackson, and many other

² "In Defense of Objectivism about Moral Obligation," *Ethics* 121 (2010), 89. Graham later modifies the subjective account by noting that most who are subjectivists also hold the view that it must be 'ability constrained' as well as evidence-sensitive. This builds into the account the intuition that 'ought' implies 'can'. He acknowledges, however, that this additional constraint is inadequately motivated.

writers who argue in favor of a subjective standard.³ However, Jackson, in his own account of a subjective standard of ‘right’ does appeal to the agent’s actual beliefs, though he idealizes the agent’s valuing: a right action on his view involves appeals to the agent’s beliefs and what the agent ought to desire at the time of action.⁴ An agent with the wrong desire set may end up viewing the wrong action as the right one. But, on his view, as long as her desires are in the right place, her action is right as long as it is what she would believe maximizes the good given her other beliefs. Technically, Jackson’s account is not evidence-sensitive in the way that Graham spells out: it is possible for someone to hold beliefs against the evidence, and yet, on Jackson’s view, still do the right thing. Her ‘evidential situation’ dictates one set of beliefs, and yet she is insensitive to her evidential situation. Perhaps self-deceptive people are like this.

In any case, depending on what problem one is primarily worried about, the subjective account can be spelled out in a variety of ways.

A standard of moral evaluation, *S*, is *evidence-sensitive* when the evaluation (according to *S*) is made relative to either the evidence the agent has at the time, of action, is available to her at the time of action, or she is aware of at the time of action.

The evidence sensitive approach may or may not involve the agent’s psychology, depending on how the evidence-sensitivity is spelled out. If the standard is spelled out so that the requisite evidence is something the agent is aware of, then psychology is brought

³ *Consequentialism*, see chapters 5 and 6.

⁴ “Decision-theoretic Consequentialism and the Nearest and Dearest Objection,” *Ethics* 101 (April 1981), 461-482.

in. If, however, the evidence-sensitivity is understood in terms of evidence the agent should be aware of (though she may not be), then it is not brought to bear in the standard. A different subjective standard more explicitly cites the agent's psychology.

A standard of moral evaluation, *S*, is *psychological state-sensitive* when the evaluation (according to *S*) is made relative to the agent's psychology, that is, the beliefs and/or desires she has (either consciously, or recoverably).

Given this type of standard, with respect to 'right', we could hold that an action is right iff the agent believes the action will produce the best results. Or, more plausibly, given what the agent believes, the action in question would produce the best results. Thus, if Melissa believes that suffering is bad, if she believes that alleviating suffering is morally good, and if she believes that the best way to alleviate suffering is to give to famine relief, then the right action (for her to perform) is to give to famine relief (even if she fails to put these all together herself).

Why make this distinction? The distinction is important in that one can isolate discrete problems for the objective approach if one understands what different people are attracted to in the subjective approach. In the case of the psychology-sensitive standard, the problem it seeks to avoid is the agent-ownership problem: on the objective view an action's rightness is independent of the agent's psychology at the time of action, and this is deemed 'alienating.' For the agent to be truly responsible, so the criticism goes, the agent needs to 'own' the action in terms of justification. The conditions of justification cannot be alien to the agent, rather, those conditions need to be present in some way in

the agent's psychology.⁵ Peter Railton tried to solve this problem by arguing that one can be committed to the objective standard without it dominating each decision one makes, that was his attempt to solve the alienation problem.⁶ The commitment to, and identification with, the right set of values as a regulative ideal is compatible with developing dispositions to act in ways that don't explicitly appeal to those values at the time of acting on the disposition. On Railton's view, this is the way that the sophisticated consequentialist proceeds. It is an elegant solution, though attacked by Frank Jackson, among others, as failing to lead to the right answers in certain cases. Another strategy – compatible with Railton's, but in principle, separable – is to simply divorce a standard of right from a standard of praise and blame and hold that the standard of praise and blame is psychology-sensitive. We will return to this suggestion.

The evidence-sensitivity standard is concerned to avoid another problem. That is the problem of moral luck in justification. A fully objective standard of right seems to be unfair in that it seems to hold people responsible for something they have no control over. If the evidence that a course of action would prove disastrous is unavailable to me, then how can I be held responsible in performing that action? Writers such as Frances Howard-Snyder and Elinor Mason are concerned that the objective standard thus, in effect, violates the most morally compelling formulation of "'ought' implies 'can'" since

⁵ This line of criticism was famously developed by Michael Stocker in "The Schizophrenia of Modern Ethical Theory," *Journal of Philosophy* 73 (1976), 453-66, and also picked up on in Jackson, *ibid*.

⁶ "Alienation, Consequentialism, and the Demands of Morality," *Philosophy & Public Affairs* 13 (1984), 134-171.

it holds that the right thing will be beyond the *practical* means of the agent.⁷ It may be logically and physically possible to do the right thing, but due to epistemic limitations, not psychologically possible.⁸ Again, this problem is addressed by the objective consequentialist by distinguishing the standard of right from the standard of praise and blame. The different subjective standards highlight different ways in which actions can go morally wrong and deviate from the objective. This indicates an important role a commitment to the objective standard plays in providing a *taxonomy* of moral mistakes for the responsible agent to be sensitive to.

It is important to keep track of this distinction. While many people hold the subjective view to be the same as the evidence-sensitive view, this isn't always the case. Donald Regan defines the objective view, in opposition to the subjective view, as one in which "...the agent's beliefs about his obligation or about the state of the world do not determine what he should do..."⁹ This is clearly a psychology-sensitive standard. On the other hand, those who opt for a 'foreseeable consequences' standard are opting for one that is evidence-sensitive.¹⁰ What is foreseeable may not be actually foreseen by the agent, but it will also not count as 'foreseeable' unless the evidence available supports it.

There is some disagreement in the consequentialism literature that centers on which way to understand, at the very least, praiseworthiness and blameworthiness – is it tied to what is recoverable from the agent's psychology, or from what the agent 'ought'

⁷ Frances Howard-Snyder, "It's the Thought that Counts," *Utilitas* 17 (2005), 265-81; Elinor Mason, "Consequentialism and the 'Ought'-implies- 'Can' Principle," *American Philosophical Quarterly* 40 (2003), 319-331.

⁸ Of course, much hinges on what is meant by 'psychologically' possible. I discuss this in *Consequentialism*, chapter 5.

⁹ *Utilitarianism and Cooperation*, 12.

¹⁰ See, for example, Bart Gruzalski's "Foreseeable Consequence Utilitarianism," *Australasian Journal of Philosophy* 59 (June 1981), 163-176.

to believe on the basis of available evidence? But even granting a distinction between ‘right’ and ‘praiseworthy’ there is an attack on the objective account of right that focuses on its *pointlessness* or failure to do anything. It is this worry I would like to address.

Carving the space can be even more complicated when we factor in the different psychological states an agent can have. For example, it isn’t merely beliefs that factor into what an agent decides to do, the agent’s desires provide the motivational impetus. If the agent’s desires are fundamentally very bad desires, if the agent wants to do things he ought not want to do, then this can render the subjective standard of right quite dubious – that is, if the standard is understood by appeal to the agent’s *actual* beliefs and desires. This is one reason Frank Jackson has articulated a standard which appeals to the agent’s actual beliefs and what the agent *ought* to desire, given the right value function. But the standard is only subjective part way – it is evidence-sensitive, but not fully psychology-sensitive.

Persons who advocate a purely psychology sensitive standard of right (or ‘ought’) hold that an agent acts rightly if she does what she believes is best, even when her understanding of ‘best’ is based on false value claims.¹¹ The intuition here is that if an agent is functioning properly as an agent, she will display the correct *formal* features in her deliberation. So, for example, if Bob believes that perfecting human nature is our proper moral goal, and that this perfection requires that we suppress all emotion, then he *ought* to work to suppress emotion; this would be the right thing for him to do, even if his view of value was completely mistaken and leads to what is, objectively speaking, a defective sort of motivation. This is because the failure to do what one believes to be

¹¹ Brian Hedden articulates such a view of ‘ought’ in “Options and the Subjective ‘Ought’,” *Philosophical Studies*, forthcoming.

best (and here I am indifferent about how ‘best’ is understood), is a failure in one’s abilities as an agent. It is a failure that denotes a severe structural flaw in agency itself, when that agency is not guided by the agent’s beliefs and values. There is something to this intuition that if someone acts in ways that don’t reflect their beliefs and desires, there is something seriously wrong with the person. Indeed, one might even wonder if the person is *acting* at all. This is too radical. A person can, loosely speaking, be weak willed and act in ways she disapproves of. But to limit wrongness, or even blameworthiness, to these sorts of cases is much *too* limiting. A person’s own critical practices have application beyond what she presently believes and desires. A person – to borrow an example from Mark Twain – may look back on her life and be ashamed of her heartless youth, judging herself to have been morally mistaken. We can’t account for that sense of moral error without appealing to a standard that is *at least* not psychology-sensitive.

But the subjectivist will concede this point, perhaps, and instead opt for one of the evidence-sensitive approaches, or even a mixed approach. If one holds that the right action is the one that the agent would perform given she is well motivated and has all available information, the standard is evidence-sensitive but not psychology sensitive. Bob is not acting rightly given this standard. But this intermediate between the fully subjective and objective strikes me as an unhappy compromise if what we are considering is the standard by which we measure moral success. Better to keep to the classic distinction between ‘right’ and ‘praiseworthy’ and note varying degrees of ‘praiseworthy’ as deviations from the standard are realized.

Again, even granting a distinction between ‘right’ and ‘praiseworthy’ there is an attack on the objective account of right that focuses on its pointlessness or failure to be of any practical use. I find the criticism odd, since the standard is intended as one of evaluation. However, I also find it false given what we are talking about is a *commitment* to the standard.¹²

As an example, the following offers a basic rendering of an objective consequentialist standard of right.

(SR) the right action is the one that produces the best *results* amongst the action options open to the agent at the time of action; it is the action that the well-motivated, fully informed agent would perform in those circumstances¹³

First, there are very many different forms the standard can take when spelled out in more detail: one can hold that best is understood in terms of what promotes objective probability of the maximally good outcome; or one can hold that best is understood in terms of what a risk averse agent would do in avoiding the worst outcome. There are many possibilities. What is key about a pure objective standard? The following: (1) right is understood independent of the agent’s actual psychology and/or (2) right is not ‘evidence-sensitive’ in the ways discussed above.

Second, it is important that this is understood as a standard, and not as a decision procedure, though a decision procedure may well be extractable. This was Peter

¹² I discuss this more fully in *Consequentialism*, particularly in chapters 5 and 6.

¹³ There is some debate on whether or not ‘full-information’ should be included in the objective standard. Indeed, I am skeptical, but that particular issue is orthogonal to the project of this paper. If not ‘full’ information, then substitute ‘all relevant’ information.

Railton's view when he argued that what counts as a good decision procedure given this standard will be an empirical issue. As a standard, it sets the success condition for **rightness**. Of course, there are serious epistemological difficulties associated with success when that is understood objectively. It may be difficult to know if and when one has actually performed an action that is 'right'. This limitation is what led W.D. Ross to note that the right action is the 'fortunate' action:

If we cannot be certain that it is right, it is our good fortune if the act we do is the right act. This consideration, does not, however, make the doing of our duty a mere matter of chance. There is a parallel here between the doing of duty and the doing of what will be to our personal advantage. We never *know* what act will in the long run be to our advantage. Yet it is certain that we are more likely in general to secure our advantage if we estimate to the best of our ability the probable tendencies of our actions in this respect than if we act on caprice.¹⁴

In my view, Ross' views here are quite instructive. Indeed, some of his insights have been clarified and expanded in the service of arguing that there is a legitimate objective standard of 'right.' Ross is implicitly, note, appealing to a distinction between doing what is 'right' and doing one's duty. There is the right action, and then there is doing the best one can in trying to get there. It is one's duty to do the best one can. In setting a

¹⁴ *The Right and the Good* (Oxford: Oxford University Press, 1930, 31-32).

success condition, then, as Railton has noted, one can extract decision procedures that one has empirical evidence will help one get closer to reaching the goal. And, further, those decision procedures themselves, whatever they happen to be, will be assessable by reference to how closely they allow us to reach success, and successfully reach success. They are evaluated in terms of relative success. Thus, the best procedure is the more accurate one. Some procedures will be better or worse in terms of meeting the standard. Since the standard is what is used to evaluate other standards of evaluation, as well as the decision procedures people employ in actual practice, the standard has *evaluative* primacy relative to those procedures.

In seeing how this works an analogy with another standard might help. Suppose that Marissa is an architect in charge of building a new theatre, and the theatre owners are striving for the most comfortable seats, given certain efficiency or cost constraints. Let us also suppose that the greatest comfort – given those constraints – is achieved by allowing for exactly 12 inches between seats in the theatre. How well the standard is met in any given case is an objective matter, and it may, in practice, be difficult to measure exactly 12 inches between seats. But whatever method is employed by the workers to measure the distance between seats, the adequacy, the goodness, of that measure is dependent on how close such a method gets one to the objective standard. The methods or procedures themselves are assessed, and assessed in terms of evidence regarding how close they get to the standard.

Other writers have drawn analogies between moral evaluation and other sorts of *normative* evaluation. So, for example, if one judges the aim of belief to be truth, then simply a person's believing that the belief she holds is true is not sufficient for epistemic

success. Of course, there are subjectivist counterparts in epistemology – these are persons who hold that epistemic success is achieved through justification alone, and then one can have more or less demanding standards of justification, as one sees in the literature on moral evaluation. In some of Roderick Chisholm’s writings he seems to adopt the view that we ought, *epistemically*, believe reasonably and avoid believing unreasonably.¹⁵ The latter stipulation is important to note. Responsibility involves avoiding mistakes as well as achieving successes in absolute terms.

On the view that I have been sketching so far, the objective standard is the standard by which we understand the normative import of yet other standards, such as that of **praiseworthiness**. Here is an example that incorporates both evidence-sensitive and psychology-sensitive elements.

(**SP**) the praiseworthy action is the action that the agent (reasonably) expects will produce the best results amongst the options she perceives to be open to her at the time of action

“Reasonably” appeals to evidence-sensitivity; “expects” and “perceives” are the psychology-sensitive elements. There are numerous ways to combine these elements in a standard of praiseworthiness. My view is the appropriateness of standards like (SP) is a function of pragmatics – what best serves utility in the long run. Thus, the objective standard of right guides both the selection of standards of praiseworthiness as well as, ultimately, the appropriateness of decision procedures for agents to follow.

¹⁵ He makes this claim in “The Place of Epistemic Justification,” *Philosophical Topics* (Spring 1986), pp. 85-92.

Many of those – at least, in the consequentialism literature – who object to (SR) believe that something like (SP) is the standard not only of praiseworthiness, but also of rightness. Extracting a plausible moral guide from such a standard seems quite intuitive. This is because success is reached via reasonable attempts. Thus, one can recommend that people try, for example, to get the best information possible relevant to the decision at hand, and then act on what they expect given that information will produce the best results. This has the supposed advantage of reducing the impact of luck on rightness. Again, in the consequentialist literature, for the objective consequentialists these two standards come apart. Note, also, that the decision procedures extracted from (SR) and (SP) could, in principle, be the same. Whether the decision procedures diverge will again be an empirical issue.

The more difficult claim to defend is the one that something like (SR) has evaluative primacy with respect to (SP). Indeed, a set of problems for the evaluative primacy of the objective has developed over the years. What unites these problems is the idea that there will be situations in which agents ought not pursue the best option; that pursuing the best option runs a significant risk of leading to the worst outcome. Thus, one *ought* to pursue a suboptimal outcome that has less, or even no risk, of leading to the worst outcome. But the key factor is that, though (SR) still provides the standard of right, there are practical worries about successfully meeting such a success condition. The practical worries, having to do in this particular case, with the epistemic limitations we often labor under, may imply that a sophisticated consequentialist not try to do what has any chance at all of the best outcome, because the chance is small relative to the chance of generating a relative disaster. This would mean that (SP) is not *the* standard of

praiseworthiness, at least, not without significant modification. One such modification would be to include consideration of avoiding disaster.

Frank Jackson presented a case such as this in his argument for a (somewhat) subjective standard of right. He builds up to the crucial case by tweaking our intuitions about standard uncertainty cases. The case involves Jill, a doctor, and her patient, John who is suffering from a non-fatal skin condition. Jill needs to decide how best to treat John:

She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition, but will not completely cure it. One of the drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug.¹⁶

Jackson goes on to note that clearly Jill should prescribe A, even though A is clearly not the best option, where that is understood as the option which maximizes the good. Thus, it looks like Jill ought to do what she knows is wrong. This is very odd, and in cases such as this can be iterated so that a policy of acting wrongly (on the objective view) can seem to be morally best as well. This case has other parallels in the literature, most notably Derek Parfit's mine-shaft case. In this example, Parfit asks us to consider:

¹⁶ "Consequentialism and the Nearest and Dearest Objection," 463.

Mine shafts: A hundred miners are trapped underground, with flood waters rising. We are rescuers on the surface who are trying to save these men. We know that all of these men are in one of two mine shafts, but we don't know which. There are three flood-gates that we could close by remote control.

The results would be these:

		The miners are in	
		Shaft A	Shaft B
We close	Gate 1	We save 100 lives	We save no lives
	Gate 2	We save no lives	We save 100 lives
	Gate 3	We save 90 lives	We save 90 lives ¹⁷

Clearly what we ought to do in the ordinary dominant sense of 'ought' is to close Gate 3, thereby doing what in the objective sense we know to be wrong, and yet also picking the option with the least overall down side. We avoid the disaster of loosing 100.

Parfit's case is based on an example presented by Donald Regan:¹⁸

		Poof	
		Push	Not-Push
Whiff	Push	10	0
	Not-push	0	10
	Override	9	9

¹⁷ *On What Matters* (Oxford: Oxford University Press, 2011), V. I, 159.

¹⁸ *Utilitarianism and Cooperation*, pp. 264-5, ft. 1.

Regan presents this as a variation on a standard Whiff/Poof case. In such cases the moral universe consists of just two persons, Whiff and Poof. The case is presented as one in which each of these persons has a button that she may or may not press, and:

Suppose that from Whiff's point of view the subjective probabilities are that Poof is equally likely to push or not-push. It is clear that subjective AU [Act Utilitarianism] requires Whiff to override (producing a subjective expected value, and indeed a certain value, of 9 units) instead of either pushing or not-pushing (each with a subjective value of 5 units), even though overriding cannot possibly be the best act in the circumstances....¹⁹

This case is much less dramatic than the doctor case and the mine-shaft case. We are not worried about anyone dying. Instead, what is at stake is a small loss of utility. And yet it still seems true that the agent 'ought' to 'override' for a guarantee of 9 units instead of a 50% chance at 10 units. The agent ought to do what the objectivist views as the wrong thing to do, it seems, even when the losses are fairly small.

In Jackson's version of the case, the judgment that Jill ought to opt for A, that this is the right thing for her to do, is a problem for the objective consequentialist because, according to Jackson, the objective consequentialist such as Railton is committed to a decision procedure, based on the standard of rightness, which involves the agent "...setting ...the goal of doing what is objectively right – the action that has in fact the

¹⁹ Regan, 265.

best consequences – and then performing the action which the empirical evidence suggests is most likely to have this property.”²⁰ I have argued elsewhere that this is a very uncharitable reading of Railton’s suggestion.²¹ The more natural way to view the objective consequentialist’s remarks on decision procedures is to separate out two different issues: the *standard* of praise and blame is one issue, and the recommended decision procedure another. What the objective consequentialist is best regarded as holding is the view that the standard of praise and blame is to be explained in terms of the standard of right – that there is an explanatory and evaluative primacy to the standard of right. Then, what decision procedures help us to arrive at the best overall outcome is to be determined based on what standard we are seeking to satisfy. But, ultimately, the evaluative standard is fully objective.

Also, the Whiff/Poof case, Jill, and the Mine-shaft case pose a supposed problem for the objective standard of right because it sounds odd to say – as the objective theorist should in these cases – that “He ought to do an action that is wrong.” But on the objective view the oddity does not translate into inconsistency. There is nothing at all incoherent in this claim. The objective theorist gives a perfectly coherent account of both why the suboptimal option is wrong, and yet why one ought, nevertheless to do it. Further, there are numerous other cases in the literature demonstrating that this oddity is not escapable by subjective consequentialists. Satisficers have argued that there are cases – at least, imaginable cases – in which there is no best outcome. One such case is

²⁰ Jackson, 467.

²¹ See my discussion in *Consequentialism* (Routledge, 2012).

suggested in the writings of Michael Slote.²² Imagine that one is a heroine in a fairy story. One has found a genie bottle, uncorked it, and the genie grants one a wish: as much money as one names! How much does one ask for? A million dollars? But wouldn't that be irrational, since one could ask for more, such as a million and one dollars? Indeed (assuming that one would be able to translate the money into some other good), where does one draw the line in picking a number to ask for? Satisficers invoke such cases to try to show that maximizing versions of consequentialism are wrong because there are cases where there is no maximum, and it is perfectly rational to (and moral) to fall somewhere short, as long as it's good enough. I am not a satisficer, but the case does plausibly demonstrate a situation in which there is no right answer so whatever one does one acts wrongly. In this case it is also true that "He ought to do the action that is wrong" whether or not one is an objectivist or a subjectivist. A subjectivist is stuck with the odd locution in any case.

Jackson himself opts for a standard that *seems* to be a mixture of the evidence-sensitive and the psychology-sensitive. The right action is a function of the agent's actual beliefs and what the agent ought to desire (or value). Given how I have specified the two distinct forms of sensitivity that seem to figure into subjective standards, Jackson's standard is psychology-sensitive in that it appeals to the agent's actual beliefs, but (perhaps) evidence-sensitive in that it also appeals to what the agent ought to value. But the latter will be an evidence-sensitive factor only if we view the value as something

²² This is not exactly Slote's case, but is similar to one he presents on p. 147 in his section of "Satisficing Consequentialism," written with Philip Pettit, *Proceedings of the Aristotelian Society, Supplementary Volumes*, 58 (1984), 139-163.

there is evidence for. Since this seems like a plausible assumption to make, I view Jackson's standard as one that mixes the two sorts of sensitivity.

Peter Graham has argued that one objectivist strategy to undermining Jackson's counterexample is to deny certain assumptions that he makes.²³ Namely, most importantly, the assumption that the morally conscientious person is solely focused on doing the right thing. Graham points out that she must also be concerned with avoidance of wrongdoing. This, then, leads to a potential conflict that occurs in cases like Jackson's. Further, in the case of wrong-doing, we can look at the seriousness of the wrong to be avoided – there will be greater moral pressure to avoid worse outcomes, for example. Thus, as Graham notes, the morally conscientious person not only tries to do what is right, but also tries to avoid doing what is wrong, and what is really, really, wrong in particular. This seems quite clear in the cases the literature has tended to focus on, Jackson's doctor case and Parfit's mine-shaft cases. The Whiff and Poof case illustrates that the avoidance of disutility is important even when the disutility is relatively small.

Thus, one way for the objective consequentialist to respond is to note that there is the best, which I cannot be confident of, and there is what seems on balance best, given my own awareness of my serious epistemic limitations. One can then keep the standard of right, but to accommodate Jackson style cases, modify the standard of praiseworthiness to something like the following:

(SP*) An action is praiseworthy if the well-motivated agent performing the action would judge it to be, on balance, best amongst her alternatives at the time of action

²³ op. cit., 95 ff.

“On balance” is then unpacked to account for *risk aversion* of the rational agent. One’s ignorance – glaringly apparent in the Jackson and Parfit style cases – is something the epistemically responsible agent should be sensitive to. This, then, ought to be factored into the standard of praiseworthiness. We need to worry about things going badly, just as we need to worry about things going well.

Someone guided by a desire to do the best he can will also use the same criteria that establish what counts as best in guiding his choice about how to act. With regard to praise and blame, risk is very important, and the responsible agent considers what happens if he *fails*. But success is determined relative to the standard, not relative to the agent’s *attempt* to meet the standard. The role of the objective standard of ‘right’ is to set a comprehensive standard for moral success by which other standards measure, or define themselves.

Thus, the competing standards point to the fact that different things go wrong in the performance of the morally right, as well as the morally praiseworthy actions, and that we are not measuring success -- even an action’s success – along a single parameter.

To see this more clearly, we can create a taxonomy of moral failure that appeals to this standard. An agent might know what to do, but not care; or might have all the information she needs to determine the best course of action, but some misguided set of values. These lead to motivational failures. The other failures are epistemic. Lack of all relevant information leads to errors; lack of information that is available to the agent leads to errors; and, most unforgivingly of all, lack of utilization of the information one actually has in one’s possession.

Taxonomy

Relevant factors	Motivation/desire	Full information Not psychology-sensitive	Available evidence Not psychology-sensitive	Evidence 'recoverable' from agent's psychology Psychology-sensitive	Evidence agent is consciously aware of Psychology-sensitive
Objective right: what the well motivated, etc. would do	X	X			
Praiseworthy 1	X		X		
Praiseworthy 2	X			X	
Praiseworthy 3	X				X

On the objective view, an agent fails to perform the right action due to numerous factors

– failure to be properly motivated, failure to have all the relevant information, failure to have the available information, and failure to properly utilize the information one has.

There is also clumsiness and ineptitude. Further, and this is what strikes many as absurd

– the failure can also be due to sheer accident, and sheer bad luck. On Jackson's view the

only way an action counts as a moral failure is when the agent lacks the appropriate

motivation and fails to utilize the information she has. If an agent intends well, and acts on what she believes (even if more or better information was available to her) she has acted rightly, and in a morally praiseworthy manner.

Why think that failures of information, will, and sheer accident are *moral* failures? They are moral in the sense that they are morally relevant. It matters to how things turn out. There is an empirical assumption being made. One is that willing the good is at least correlated with good outcomes; that information improves the level of control we exert over the world, and thus over the effects of what we do. These are entirely reasonable assumptions, as is the following:

(MRel) If factor *f* is a factor that one ought to take into consideration in performing an action, then it is morally relevant

and

If a factor is morally relevant then the objective moral standard is sensitive to it

Consider the appropriate response to such a failure: imagine that Jackson's doctor, Jill, failed to attend a vital seminar – to which she was invited – that would have provided information about the skin condition that would have enabled her to have treated the patient more effectively. That she lacks such information does reflect badly on her. What guides the standard of praise brought to bear in a given case reflects our justifiable

expectations regarding how a person should be held responsible in the service of promoting the good. The overarching standard is thus the objective standard of 'right'. What makes it a mistake – a morally significant mistake – that the doctor did not go the seminar? If she had gone the outcome would likely have been much better for the patient. What is it to be *better for* the patient? Closer to the best, *really* best, outcome for the patient. The aim of morality, including practical morality, is to achieve the good, not to achieve what one just happens to believe to be good.

It may be that what is behind the view that certain mistakes, while making achievement of the good less likely, are nevertheless not morally relevant, is the view that the well-motivated agent, the agent whose *desires* are guided by the right sorts of value commitments, has fulfilled whatever is demanded of her morally already, and that the content of her (non-moral) beliefs is not something that can affect that. On this view, however, an agent can do no wrong no matter how ill-informed, as long as she wants to act rightly. If this is understood Jackson's way, as being only partially psychology-sensitive – then it is much more plausible than the purely subjective view. However, it still seems quite implausible in that if such a standard is the standard of right that is primary – against which the others are compared – then there is no room for moral improvement. Yet arriving at better outcomes, and taking pains to acquire information that makes sure one arrives at better outcomes, does constitute moral improvement. Thus, the purely objective standard is the primary standard.