

Objective Morality, Subjective Morality, and the Explanatory Question[†]

Dale Dorsey

Department of Philosophy
University of Kansas
1445 Jayhawk Boulevard
Wescoe Hall, rm. 3090
Lawrence, KS 66045
ddorsey@ku.edu

David Hume’s assessment of John Hambden, one of the principal instigators of the English Civil War, runs as follows:

The attempt. . . of totally annihilating monarchical power, was a very blameable extreme; especially as it was attended with the danger, to say the least, of a civil war, which, besides the numberless ills inseparable from it, exposed liberty to much greater perils than it could have incurred under the now limited authority of the king. But as these points could not be supposed so clear during the time as they are, or may be, at present; there are great reasons of alleviation for men who were heated by the controversy, or engaged in the action.¹

For Hume, there are two ways of assessing the conduct of Hambden. First, one could consider whether the action was genuinely right or wrong: whether, taking in all the facts, Hambden conformed to his moral obligations. In this light, Hume contends, Hambden’s action was a “very blameable extreme”. But Hume also investigates Hambden’s moral propriety in light of the information available to Hambden at the time. From this perspective—if Hume is correct that the risks could not have been fully foreseen or understood and that Hambden’s heart was in the right place (especially when it comes to his concern for liberty)—there do seem to be “great reasons of alleviation”, or reasons to temper our otherwise harsh assessment of Hambden and his co-conspirators.

[†]Thanks to Janice Dowell, Clayton Littlejohn, and a number of anonymous reviewers.

¹David Hume, *The History of England* v. 5 (Indianapolis, IN: Liberty Fund, 1983 [1778]), 574. I want to thank [...] for calling my attention to Hume’s treatment of Hambden. Whether my reading of Hume’s treatment is faithful to Hume’s own position is, for my purposes, neither here nor there.

I find these thoughts attractive, even platitudinous. Indeed, a common presupposition in metaethical theory is that moral assessment comes in (at least) two flavors, one of which is more-or-less sensitive to our epistemic circumstances, the second of which is not so sensitive.² On this view, one can have non-coextensive *objective* and *subjective* moral obligations. Hambden’s action, at least according to Hume, is *objectively* wrong—a “blameable extreme”—but *subjectively* right: morally appropriate in light of Hambden’s beliefs.

Though these thoughts are attractive, a number of questions arise. In this paper, I limit my discussion to what I dub “the explanatory question”: how one might understand, in a more or less precise way, the construction of one’s subjective moral obligations *given* one’s objective moral obligations. I argue that a proper answer to the explanatory question is important not simply for its own sake, but that it also sheds new light on important challenges to the existence of both objective and subjective moral obligations. In this paper, I do not seek to defend the “two oughts” thesis in any robust way. Rather, I intend to ferret out its best answer to the explanatory question, and the importance of my preferred answer for wider debates concerning the overall tenability of a distinction in objective and subjective morality. The plan of the paper runs as follows. In §1, I introduce the distinction between objective and subjective morality, its standard attractions, and the explanatory question. In §2 I characterize, and offer a critical analysis of, the traditional answer to the explanatory question. In §§3-5, I attempt to provide a somewhat more precise characterization of the explanatory relationship between subjective and objective moral obligations (or ‘ought’s). Beginning in §6, I argue that my conception of this relationship can blunt the force of a few notable objections to the importance of both objective and subjective morality.

1. Preliminaries

To explain in more detail the explanatory question, a few words must be said about the general view that gives rise to it. As I understand it, this view is characterized by the following rough tenets:

1. Morality permits of two distinct perspectives of assessment.
2. A first, “objective”, perspective assesses the moral quality of actions in a way that is not sensitive to agents’ epistemic cir-

²See Henry Sidgwick, *The Methods of Ethics* (Indianapolis, IN: Hackett Publishing Company, 7th ed., 1981 [1907]), 394-5.

cumstances.

3. A second, “subjective”, perspective assesses the moral quality of actions in a way that is sensitive to agents’ epistemic circumstances.
4. Objective moral assessment is explanatorily prior to subjective moral assessment.³

For lack of a better moniker, call this collection of theses “Objective and Subjective Morality” (OSM).

OSM subscribes to the general thesis that there are two independent moral perspectives, along with two distinct moral ‘ought’s that correspond to each perspective. Notably, however, OSM sets a priority rule: subjective morality is subsidiary to, or explained in terms of, objective morality. One might put this priority rule in a slightly different way. If subjective morality is explained in terms of objective morality, then it would seem that the fundamental moral facts, i.e., the fundamental facts that determine whether we conform to our moral obligations or not are *objective*, i.e., independent of the epistemic circumstances of individual agents. Subjective moral obligations are defined in terms of the more fundamental objective facts.

Why posit an independent subjective morality, if we assume that the fundamental moral facts are objective? Why, in other words, care about whether Hamden made the right decision, if we already know that the fundamental moral facts—objective moral facts—declare that he did the wrong thing? I can think of two reasons, one direct, one indirect. In conforming to our subjective obligations, though this does not guarantee that we will do the right thing in all cases, it does say something important about our *character*. Performing the subjectively right act seems to indicate that we are disposed, or oriented, to conform to our objective moral obligations. A further way to put this point is that conforming to one’s subjective moral obligations is evidence of a virtuous form of moral decisionmaking: one decides to act in ways that, at least in light of one’s epistemic circumstances, are the right way to act.⁴ This reason is indirect in the sense that its

³Notice that I refrain from identifying OSM as committed to the claim that the objective ‘ought’ is explanatorily prior to the subjective ‘ought’. Though this is one popular way to understand OSM, it isn’t the only way and, so I argue in the next section, it is false.

⁴Of course, this evidence is imperfect for two reasons. First, I could engage in virtuous moral decisionmaking at t_1 , drop dead at t_2 , and hence not conform to my subjective moral obligations at t_3 , as I (virtuously) decided to do at t_1 . Furthermore, I could conform to my subjective moral obligations without explicitly engaging in any sort of virtuous decisionmaking. I could, for instance, simply conform to such an obligation by

importance is derived from the importance of conforming to the explanatorily prior objective ‘ought’. However, there is a further, direct, reason to care about the subjective rightness of our actions: whether our actions are subjectively right seems intimately tied to the appropriateness of reactive attitudes, punishment, or blame. In Hambden’s case, given that his action was subjectively right, we are less likely to treat blame as appropriate (we would admit “great reasons of alleviation”).⁵

The first consideration in favor of OSM is somewhat technical. Accepting a distinction between objective and subjective moral requirements seems to solve a number of paradoxes, including those involving practical conditionals. To see this, take the following case, from Kolodny and MacFarlane:

Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed. We take it as obvious that the outcome of our deliberation should be

(1) We ought to block neither shaft.

Still, in deliberating about what to do, it seems natural to accept:

(2) If the miners are in shaft A, we ought to block shaft A.

(3) If the miners are in shaft B, we ought to block shaft B.

We also accept:

(4) Either the miners are in shaft A or they are in shaft B.

But (2), (3), and (4) seem to entail

(5) Either we ought to block shaft A or we ought to block shaft B.

And this is incompatible with (1). So we have a paradox.⁶

luck. Though this evidence is imperfect, however, there remains good reason to take it, rather than conformity to the objective perspective, seriously in assessing an individual’s decisionmaking. Only this perspective offers any evidence concerning the quality of a person’s actions in light of their beliefs.

⁵It is worth noting that different moral theories will interpret the importance of conformity to the subjective and objective moral ‘ought’s differently. For instance, act-consequentialism, in doling out punishment, will treat the moral quality of one’s decisionmaking as important, but only insofar as punishing those who make poor moral decisions produces best consequences (as it will in most cases).

⁶Niko Kolodny and John MacFarlane, “Ifs and Oughts” in *Journal of Philosophy* 107 (2010).

One might be tempted to reject one of (1)-(3). One might be (as Kolodny and MacFarlane are) tempted to reject *modus ponens* for ‘ought’ claims of this kind. But if we accept that ‘ought’ comes in two forms, we can avoid paradox in a particularly natural way. The ‘ought’ in (1) refers not to the genuine moral valence of our action, but rather to the best *decision*, given the information we have available. Hence the ‘ought’ in (1) is subjective. The ‘ought’s in (2)-(5) are objective: they are meant to indicate that either blocking shaft A or blocking shaft B is morally permissible depending on whether the miners are actually in shafts A or B. Hence (5) is not incompatible with (1). We can surely make the right decision (to block neither shaft) while performing the wrong act (failing to save all the miners).⁷

To solve the above paradox, of course, it is sufficient to simply make a distinction between objective and subjective moral ‘ought’s. But it is not enough to defend the claim that the former should be explanatorily prior to the latter. However, one might also consider the following:

Roger’s Murder: Dave is on trial for Roger’s murder. Many trustworthy witnesses put Dave at the scene of the crime, and there is very strong physical evidence suggesting that Dave, in fact, killed Roger. But, as it happens, Dave was framed, rather convincingly, by Miles.

Suppose that it is important, from a moral perspective, to punish those guilty of murder. Given trustworthy testimony from a wide array of witnesses, along with the physical evidence of the case, it is perfectly morally appropriate to decide to punish Dave. But if one learns at a later date that Miles, in fact, framed Dave, (perhaps Miles confesses in a late-in-life bid for religious salvation) what would our reaction be? Under such circumstances, we would make a distinction between the moral propriety of the act itself, and the moral propriety of the action given the information we had (which would be evidence of, in turn, the moral propriety of our moral decisionmaking, etc.). Our untutored reaction is likely to be that though we made the right decision in light of our evidence (i.e., our action is subjectively right), we did the wrong thing (i.e., our action is objectively wrong, or wrong in light of all the facts). Furthermore, given that punishing Dave was subjectively, if not objectively, right we would be unlikely to subject ourselves to blame, nor would we regard our moral character as tarnished, as a result of acting in a morally prohibited way. However, it seems to me that a significant *reason* we would not be appropriate targets of moral blame insofar

⁷Note that this solution to the problem is subject to a number of responses, only one of which I will address here (§6.2).

as we conform to our subjective obligations is in part explained by OSM's explanatory priority. The reason blame is not appropriate in this case is that we were "oriented" in the right way: oriented to that which would be objectively right were our beliefs correct. And if conforming to one's subjective obligations entails that one is properly "oriented" toward that which is objectively right, it follows that subjective obligations should be explained, at least in part, by objective obligations.

Some will contend that OSM is very closely tied to what might be called "objective act-consequentialism", or the view that the moral valence of actions is determined by the relative quality of the actual consequences of those actions. And, indeed, those who accept objective act-consequentialism are likely to accept something like OSM (insofar as they accept a subjective perspective at all). However, OSM is not confined to act-consequentialism. Any moral theory that treats propositions *about which we may be mistaken* (such as, e.g., "Dave is guilty") as genuine moral reasons can sensibly make a distinction between the way in which the morality of an action appears to us in light of what we believe and the genuine morality of that action. Accepting OSM will allow, then, that the objective assessment of actions determines moral valence but that subjective rightness can assist in a determination of the warrant of praise or blame and a determination of individual moral character.

However, for whatever moral theory we accept, OSM remains ambiguous in one important way. Though OSM declares that objective morality is explanatorily prior, we appear to lack an account of the method by which the objective perspective explains the subjective. Given that I objectively ought to ϕ , how do we determine that which I subjectively ought to do? Call this "the explanatory question". This paper is designed to critique the traditional answer to the explanatory question, and to answer it in an independently plausible way. In addition, I show that a proper answer to the explanatory question can go some distance toward alleviating some—though surely not all—skeptical worries that surround OSM.

2. *Against the Traditional View*

When it comes to the explanatory question, one answer looms large. Call this "the traditional view". The traditional view holds that one's subjective obligations are the result of a counterfactual: one subjectively ought to ϕ if and only if, on the assumption that one's beliefs⁸ are true, one objectively

⁸I use "beliefs" as shorthand here for whatever relevant set of epistemic circumstances

ought to ϕ .

This conception is expressed in the following terms by, e.g., Kolodny and MacFarlane: “ S [subjectively] ought (at t) to ϕ iff ϕ ing is the best choice available to S in light of what S knows at t ,”⁹ Zimmerman: “[a]n agent [subjectively] ought to perform an act if and only if he believes that it is the best option that he has,”¹⁰ and Peter Railton, who writes, describing “subjective consequentialism” that: “*Subjective consequentialism* is the view that whenever one faces a choice of actions, one should attempt to determine which act of those available would most promote the good, and then try to act accordingly.”¹¹ On Railton’s view, one subjectively ought to do that which in one’s best estimation will actually produce the best consequences, and will hence be the subject of an objective ‘ought’. Though these views are obviously not equivalent, they all accept the traditional view: that one’s subjective obligations are determined by a crucial counterfactual, i.e., that which, given one’s beliefs, is objectively required.

My statement of the traditional view requires one bit of disambiguation.¹² The ambiguity concerns the term “beliefs”. Some versions of the traditional view, such as Zimmerman’s, will allow that the subjective ‘ought’ is relative not just to a person’s non-moral beliefs, but also relative to a person’s *moral* beliefs, i.e., beliefs concerning the moral facts. Others, such as Kolodny and MacFarlane’s, seem to hold that an individual’s subjective ‘ought’ is sensitive only to an individual’s non-moral beliefs.¹³ For now I want to abstract from this particular ambiguity. In critiquing the traditional view, I will focus on the shared feature of both interpretations, viz., that an individual’s subjective moral obligations should be in part determined by what they objectively ought to do on the assumption that their non-moral beliefs are true.

This ambiguity notwithstanding, the traditional view is popular. Its popularity is certainly understandable. After all, if the subjective moral

one might prefer, i.e., beliefs, evidence available, evidence available in light of some further idealized perspective, etc.

⁹Kolodny and MacFarlane, 116.

¹⁰Michael Zimmerman, *Living With Uncertainty* (Cambridge: Cambridge University Press, 2008), 5.

¹¹Railton, “Alienation, Consequentialism, and the Demands of Morality” in *Facts, Values, and Norms* (Cambridge: Cambridge University Press, 2003), 165.

¹²Thanks to [...] for this helpful suggestion.

¹³I suppose there is a third possibility, i.e., that the subjective ‘ought’ is sensitive only to a person’s moral beliefs, not moral and non-moral beliefs. But this seems unmotivated, and is certainly unresponsive to the motivation for OSM as noted in the cases of Hambden, Roger’s murder, etc.

perspective is the primary means by which we evaluate decisionmaking or moral character, and if our decisionmaking is appropriate only insofar as it is directed toward the performance of actions of positive moral valence, it makes perfect sense to say that the subjective ‘ought’ is given simply by what would be objectively required in light of our epistemic circumstances. And though this view has a natural motivation, it is unsatisfactory. To see why, consider the following principle:

Explanatory Principle One (EP1): objective moral requirements are explained in terms of objective moral reasons; subjective moral requirements are explained in terms of subjective moral reasons.

EP1 holds that objective and subjective moral obligations are explained by *reasons*: in particular, objective and subjective moral reasons, respectively. The argument for EP1 runs as follows. It is surely true that objective moral requirements are determined by objective moral reasons and the respective weight of the objective reasons involved. Assume, for instance, that Miles is guilty of Roger’s murder. I can now choose to imprison Miles or hang Miles. Assume that hanging Miles will provide a modicum of vengeful pleasure to Roger’s family. This is surely a reason to hang Miles rather than to imprison him, and if this is all there were to it, I should surely hang Miles. But also assume that prison is more humane than the gallows. Given that it seems plausible to declare that the humaneness of punishment is a more significant moral factor than the production of vengeful pleasure, it would seem that the objective reason against hanging Miles is stronger than the objective reason in favor of hanging him. If so, it would seem that I objectively ought to imprison Miles rather than to hang him. This objective requirement is clearly a product of the objective reasons involved, viz., that Miles is guilty, and that though the promotion of pleasure for Roger’s family is *a* reason to hang Miles, prison’s greater humanity outweighs this reason. If so, the objective ‘ought’ seems clearly determined by objective reasons and their relative weight. (For the purposes of this paper, I will assume that moral rationality is “optimizing”, i.e., that one morally ought to perform the act for which there is the strongest balance of moral reasons. Nothing hangs on this, however; my arguments will work *mutatis mutandis* for a form of moral “satisficing” rationality,¹⁴ or any other form of moral rationality one might choose to accept.¹⁵)

¹⁴See, for instance, Michael Slote, *Beyond Optimizing* (Cambridge, MA: Harvard University Press, 1989).

¹⁵For a non-standard account of moral rationality, see Douglas Portmore, “Are Moral

But assume now that there are *subjective* reasons, and that like objective reasons, these subjective reasons have variable significance or weight. To hold that the subjective ‘ought’ is not determined by subjective reasons, one must hold that it is possible for it to be the case that the weightiest subjective reason directs one to ϕ , but that one subjectively ought to $\neg\phi$ (only if the subjective ‘ought’ is explained at least in part in terms of subjective reasons can one guarantee that one will be subjectively required to perform an action for which there is greatest subjective reason). But (a) this seems independently absurd, and (b) this seems to draw a sharp, and unwarranted, distinction between the rational structure of subjective and objective morality: one must assume that one objectively ought to perform the act for which there is the greatest balance of objective moral reasons, but that it’s not the case that one is subjectively required to perform the action for which there is the greatest balance of subjective moral reasons. That there should be such a radical cleavage between moral perspectives is unmotivated.

Hence *if* there are subjective reasons, it would seem that subjective moral obligations are explained in terms of them. But the antecedent should be affirmed. This can be seen by considering a case similar to one offered by Allan Gibbard. Imagine that, in driving to meeting, I approach an intersection at which there is a pedestrian crosswalk. Assume that I’m running late for the meeting, and that it would be a serious inconvenience for several people for me to be late (the later I am, the more inconvenienced they will be). Assume also that, though I’m not really very sure, I *believe* that there is no pedestrian coming. From this, it seems to me, two things follow. First, there is surely some subjective moral pressure to proceed at speed. After all, if I don’t proceed at speed, people will be inconvenienced, and I do, in fact, believe that there is no one coming. But, as Gibbard writes: “At a blind intersection. . . speeding on through is subjectively wrong, even if most likely no cross-traffic is coming.”¹⁶ In this case, given one’s epistemic circumstances, one subjectively ought to slow down. But if I subjectively ought to slow down even when there is subjective reason to proceed at speed, my subjective ‘ought’ can only be explained by the presence, and greater weight, of a different subjective moral reason: a reason to slow down. If so, it would appear that subjective reasons exist. And if this is correct, EP1 holds not just for the objective ‘ought’, but for the subjective ‘ought’ as well.

Reasons Morally Overriding?” in *Ethical Theory and Moral Practice* 11 (2008).

¹⁶Allan Gibbard, “Truth and Correct Belief” in *Philosophical Issues* 15 (2005), 345.

So what does all this say about the traditional view? If we accept EP1, we should be wary of the claim that objective moral requirements are explanatorily prior to subjective moral requirements. It would appear that subjective moral requirements are not explained in terms of objective moral requirements at all, but rather in terms of subjective moral *reasons*. Of course, one way to accept EP1 and to retain the traditional view might be to hold that subjective moral reasons are explained in terms of objective moral requirements.¹⁷ If this is correct, the traditional view follows: if one has subjective reason only to perform that action that is objectively rational in light of one's non-moral beliefs, it follows trivially that one subjectively ought to perform that action that would be objectively obligated in light of one's non-moral beliefs. But Gibbard's example shows that this suggestion also fails. If the strongest subjective reason I have is to slow down, it would appear that I can have a subjective reason, indeed a very strong one, to act in a way contrary to that which is objectively obligated in light of one's non-moral beliefs. After all, I *believe* no one is coming. In light of this belief, I objectively ought to proceed at speed. But I subjectively ought to slow down. Hence the traditional view fails: I am subjectively required not to perform an action that is objectively required in light of my beliefs.

One further suggestion is worth consideration here.¹⁸ Some have held that though the traditional view is false, the subjective rightness of actions is not determined by subjective reasons, but rather, simply, by the expected consequences of alternative actions available.¹⁹ This seems to adequately respond to Gibbard's case: because it is too risky, I subjectively ought to slow down rather than to proceed at speed. Hence there is no reason to believe that the subjective 'ought' is determined by subjective reasons. However, this proposal is mistaken. First, to suggest that one subjectively ought to perform the action with greatest expected utility is incompatible with a number of non-consequentialist moral theories that might also accept OSM. For instance, it may be that refusing to punish Dave has highest expected

¹⁷This approach is suggested by David Brink. Brink writes that "we identify an agent's subjective reasons with the actions that would be objectively rational if only her beliefs about her situation, or the beliefs about her situation that it would be reasonable for her to hold, were true," ("Prudence and Authenticity: Intrapersonal Conflicts of Value" in *The Philosophical Review* 112 (2003), 220). Here Brink admits the existence of subjective reasons, but explains subjective reasons only in terms of what would be objectively rational, i.e., that one objectively ought to perform.

¹⁸Thanks to [...] for suggesting that I consider this view.

¹⁹Cf. J. J. C. Smart, "An Outline of a System of Utilitarian Ethics" in Smart and Williams *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1974), ch. 7.

utility, but that given the overwhelming moral importance of punishing the guilty, one subjectively ought to punish Dave. Here it is clear that both objective and subjective reasons determine their respective ‘oughts’. However, one might alter this proposal in the following way: subjective ‘oughts’ are determined by the greatest expected *moral value*. But the best understanding of this theory of subjective obligations is not incompatible with EP1 (or, indeed, my considered view; see note 29). Even if it is the case that ϕ -ing has greatest expected moral value, we wouldn’t wish to deny that the fact that ψ -ing has positive (though non-optimal) expected moral value is a reason to ψ . If so, we should understand such a view as holding that ψ -ing has positive expected moral value constitutes a reason to ψ ; but this reason could be outweighed by weightier reasons—understood in terms of expected moral value—in favor of ϕ -ing.²⁰

3. An Alternative View: Part 1

To reject the traditional view may seem to put OSM in jeopardy. If conformity to the subjective ‘ought’ characterizes the quality of an individual’s moral character, and the subjective ‘ought’ is not determined by the objective ‘ought’, then it would seem that either one’s moral character is determined in a way that has nothing to do with the genuine moral valence of one’s actions—which seems independently implausible—or it must be that the objective ‘ought’ does not determine the genuine moral valence of one’s actions—which is a rejection of the priority rule embraced by OSM.

But I think there is an answer to the explanatory question that avoids the problems of the traditional view and can preserve OSM’s priority rule. Rather than placing the explanatory priority of objective morality at the level of obligations, we should instead place this priority at the level of *reasons*: subjective moral reasons, which in turn explain the subjective moral ‘ought’, are determined by objective moral reasons. Insofar as objective moral reasons determine the objective ‘ought’ (as per EP1), the person oriented toward objective moral reasons will maintain an orientation that plausibly reflects a good moral character: this person is oriented toward that which fundamentally determines the moral valence of one’s actions.

Two questions immediately arise. First, *what distinguishes subjective*

²⁰Note also that even if this theory is correct, it is not a response to the explanatory question: it does not explain the relationship between objective and subjective moral obligations. At most, it offers a theory—possibly true, possibly false—of what our subjective obligations *are*. Furthermore, this view is not entirely ecumenical between substantive views of one’s subjective obligations. See, for instance, note 27.

and objective reasons? Second, *how are subjective reasons determined in terms of objective reasons?* On the first question, considerable headway is made by Mark Schroeder. Schroeder writes:

Objective normative reasons, then, depend on how things are independently of the agent’s beliefs (even when the relevant way things are is a matter of what the agent believes). Subjective normative reasons depend on what the agent believes, independently of how things actually are. . . Subjective reasons are to be understood in terms of objective reasons. Bernie counts as having a reason to take a sip by virtue of believing that his glass contains gin and tonic, on this view, because the content of his belief is the right kind of thing to be an objective reason for Bernie to take a sip—were it true:

Subjective: For R to be a subjective reason for X to do A is for X to believe R , and for it to be the case that R is the kind of thing, if true, to be an objective reason for X to do A .²¹

Schroeder’s account of the distinction and relationship between objective and subjective reasons is compelling. But there is one minor problem. Take Gibbard’s case. In this case, I don’t believe that there is a pedestrian coming (though I cannot, for sure, rule it out); in fact, I believe that there is no pedestrian coming. But even if I don’t believe that there is a pedestrian coming, I still seem to have a reason to slow down *given* that I can’t rule it out. Insofar as Schroeder’s account holds that any subjective reason must be believed, any proposition I fail to believe cannot constitute a subjective reason, and hence I couldn’t have subjective reason to slow down on Schroeder’s view.²²

Leaving aside this problem, however, I prefer an account of subjective reasons identified not in terms of the *believed* proposition, but rather in terms of the belief *itself*.²³ So I want to tweak Schroeder’s view just slightly

²¹Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press, 2007), 13-14.

²²Also, Schroeder’s view also implies that that false propositions can be subjective reasons. Schroeder himself accepts this conclusion (Schroeder, 14n23), so I refrain from treating this as an objection to his view here. For what it’s worth, my alternative will avoid this, and will instead hold that a subjective reason is a proposition of the form “I believe r ”, rather than r , which is marked as subjective *because* it is believed.

²³See Mark Schroeder, “Having Reasons” in *Philosophical Studies* 139 (2008), 60-61.

(this tweak will also address the problem I discuss in the previous paragraph). Insofar as I seek to offer an account that treats objective reasons as explanatorily prior to subjective reasons, I will take, as Schroeder’s account does, objective reasons as primitive. With this in mind, subjective reasons can be defined as follows:

*Subjective**: A proposition r , if true, is a subjective reason for x to ϕ if and only if r contains a corresponding objective reason for x to ϕ within an operator that refers to x ’s epistemic circumstances.²⁴

*Subjective** holds that a subjective reason is a proposition that contains a “corresponding objective reason” within an epistemic operator (such as belief, disbelief, indifference, etc.). A corresponding objective reason is not a new type of reason, standing alongside objective and subjective reasons, but is simply the sort of thing that would count, objectively speaking, in favor of a particular action were it true. “Dave murdered Roger” is, if true, an objective reason for me to punish Dave. But “I *believe* Dave murdered Roger” is a *subjective* reason for me to punish Dave: the corresponding objective reason “Dave murdered Roger” has a positive epistemic status for me, which renders it the case that I have a subjective reason to punish Dave.²⁵

One note before I move on. Following Schroeder (and following my practice so far), I interpret the proper “epistemic circumstances” to which our subjective reasons refer as a person’s *beliefs*. I adopt this proposal merely for ease of exposition. Nothing in what follows will ride on this—readers are free to select whatever epistemic circumstances they feel most appropriate (for instance, one could say that subjective reasons are those corresponding objective reasons in which one has *evidence*, perhaps under idealized circumstances).

4. An Alternative View: Part 2

²⁴I hold that true propositions can be reasons; others will deny this. Some will insist that facts, and not propositions, are reasons. However, *Subjective** can be translated *mutatis mutandis*. Using the notion of a credence level explained below, one might say that the fact that I grant a particular credence level in p is a subjective reason for x to ϕ if and only if p refers to an objective reason for x to ϕ . Thanks to [...].

²⁵“I believe Dave murdered Roger” may also be an *objective* reason for the performance of certain actions. For instance, if it is overwhelmingly clear that Dave did not murder Roger, “I believe Dave murdered Roger” would be an objective reason for me to, for instance, investigate whether I have some sort of prejudice against Dave, or insensitivity to evidence, etc.

So far we have accounted for the nature of subjective reasons. How, then, are these subjective reasons explained in terms of objective reasons?

This question breaks down into two others. First: what is the explanandum? Second: what is the explanans? In answering the first question, it is important to note that at any given time, *many* subjective reasons will apply to me. Every proposition (and hence every objective reason) will maintain *some* epistemic status or other for me (and hence will count as a corresponding objective reason), even if *most* of the time, this status is epistemic indifference. But, given *Subjective**, even indifference toward a given proposition can generate subjective reasons in favor of a given action. Thus what must be explained is not *which* subjective reasons apply to me, but rather the relative moral significance, or *weight* of the subjective reason-constitutive propositions. For ease of exposition I will assume that the moral weight of any particular reason can be identified numerically: a weight of -1 entails that p is a reason of strongest possible weight *against* ϕ -ing, 0 means that p is neither reason for nor against ϕ -ing, 1 entails that p is a reason of strongest possible weight in favor of ϕ -ing. Thus the explanandum: a proper explanation of subjective reasons in terms of objective reasons must assign a numerical weight to all propositions that count as subjective reasons. Now to the explanans.

An example may help. Let's say I have to decide whether to punish Dave for Roger's murder. I believe that he did it, but my belief is relatively weak. The only information I have to go on is that a stool pigeon, whom I know to be correct 51% of the time, fingered Dave. I surely have a subjective reason to punish him. But how strong is this subjective reason? Surely one factor in determining the weight of this reason is the relative moral importance of punishing Dave, assuming he is guilty. To put this in terms indicated in the previous section, it is important to determine the objective weight of the *corresponding objective reason* (i.e., the objective weight of "Dave is guilty" as a reason to punish Dave). But this must be balanced by a further factor in the determination of the strength of my subjective reasons: the strength of my *belief* in this corresponding objective reason. If, for instance, to punish the guilty is of decisive moral importance, though my belief that he is guilty is weak, I nevertheless have a subjective reason to punish him: after all, to punish him on the assumption that he is guilty of the gravest moral importance. But, for instance, if it is more morally important not to punish the innocent than to punish the guilty, one might think that *because* my belief in Dave's guilt is so weak, though I still have a subjective reason to punish him, it is of insufficient strength to render it

the case that I subjectively *ought* to punish him.

To capture these general thoughts in a more precise manner, consider the notion of a “credence value” or “credence level”. The credence level of a given proposition r for a given individual A is determined by the strength of A’s belief in r .²⁶ A credence level of 1 indicates that A maintains the strongest possible belief in r . A credence level of 0 means that A grants r no credence whatever: A maintains the strongest possible *disbelief* of r . A credence level of .5 indicates that A is epistemically indifferent to r .²⁷

In light of this terminology, it seems sensible to determine the weight of a subjective reason as a function of the moral significance of the corresponding objective reason (i.e., the moral significance of “Dave is guilty” in providing a reason to punish Dave), and of my credence level *in* this corresponding objective reason (i.e., how strongly I believe that Dave is guilty). Consider now:

Explanatory Principle Two (EP2): The weight of a subjective reason sr to ϕ ($W_{sr,\phi}$) is a function (f) of one’s credence value in the corresponding objective reason r (CRE_r) and the objective weight of r in providing a reason to ϕ ($OW_{r,\phi}$): $W_{sr,\phi}=f(CRE_r, OW_{r,\phi})$.

Different substantive theories of one’s subjective moral obligations will offer different accounts of f .²⁸ A natural possibility, however, is multiplication.²⁹ Take an easy case. If we assume that the objective weight of “Dave is guilty”

²⁶Those who would prefer an alternative to “belief” (i.e., evidence, or some other epistemic stance) are free to interpret a credence level as relative to whatever epistemic stance they deem relevant.

²⁷One might put this in terms of one’s assessment of the objective chance of a proposition: a credence level of zero in a particular proposition means that you grant that proposition no objective chance of being true. A credence level of .5 means that you grant this proposition an equal chance of being true or not true. A credence level of 1 means that you grant this proposition no chance of being false. See, for instance, David Lewis, “A Subjectivist’s Guide to Objective Chance” in *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 2000).

²⁸One desideratum seems sensible. If it’s the case that if either CRE_r or $OW_{r,\phi}$ are zero, $W_{sr,\phi}$ should also be zero. This reflects the commonsense assumption that if a particular reason is of no weight in providing a reason to ϕ , it cannot provide a subjective reason to ϕ , despite its being believed. Also, if a particular reason has a credence of zero, despite its heavy objective weight, this should also provide no subjective reason.

²⁹One alternative might be to decrease the relative importance of the credence score as the objective weight increases, denying that f is straightforward multiplication; one might accept this position if one believes that even a very small chance, of, say, killing someone is a strong subjective reason not to perform a given act; I leave this suggestion aside here.

as a reason to punish Dave is 1 (given that one is absolutely required to punish the guilty), and one’s credence level in “Dave is guilty” is .51 (given the unreliability of one’s informant), the subjective reason of “I believe Dave is guilty” (*sr1*): $SW_{sr1,PunishDave} = (.51 * 1) = .51$.³⁰ Compare this to the weight of “I believe Dave is innocent” (*sr2*) as subjective reason not to punish Dave. Assuming that one believes that “Dave is innocent” with strength .49, and that the objective weight of “Dave is innocent” in providing reason not to punish Dave is 1: $W_{sr2,Don'tPunishDave} = (.49 * 1) = .49$. The subjective reason to punish him is stronger than the subjective reason not to do so, and hence there is greater subjective reason, though only slightly greater, to punish him rather than not. If so, we subjectively ought to punish Dave.

Take, now, Gibbard’s case. Assume that “there is a pedestrian coming” is a corresponding objective reason of strength .8 against proceeding at speed (given that proceeding at speed under such circumstances will result in an injured pedestrian). One might also hold that “there is no pedestrian coming” is a corresponding objective reason of strength .2 to proceed at speed (given relative moral importance of inconvenience and serious injury to a pedestrian). Assume that the respective credence values are .3 and .7, respectively. If so, the weight of “I believe someone is coming” (*sr1*) and “I believe there is no one coming is coming” (*sr2*) is: $W_{sr1,SlowDown} = (.3 * .8) = .24$; $SW_{sr2,ProceedatSpeed} = (.7 * .2) = .14$, respectively. If this is correct, there is a stronger subjective reason to slow down and hence one subjectively ought to slow down, though one believes that there is no one coming and hence that one objectively ought to proceed at speed. These seem like the right answers.³¹

Of course, if we change the corresponding values, the strength of the reasons in question may change. It could be, for instance, that the relative credence value in the claim that there is no one coming is .9 rather than .7,

³⁰By the way, the numerical values I attach to objective weights are obviously arbitrary; I would expect the true objective theory of morality to specify these values in more consistent detail.

³¹Indeed, my view is compatible with the suggestion that one subjectively ought to maximize expected value: expected value is assigned by subjective probabilities, which are expressed as credence values (see note 20) multiplied by the value of resulting states. One could express the formula given in EP2 to deliver this result: if the only objective reasons are given by propositions concerning the value of consequences of a given action, subjective reasons (if *f* is multiplication), will be entirely determined by the expected consequences of a given action, and hence one will be subjectively required to maximize expected value. Of course, we are not compelled to adopt this approach: we could accept a wider range of moral reasons, or we could reject the claim that *f* is multiplication. But my view is compatible with this standard proposal.

say. If that's right, there may be subjective reason enough to proceed at speed. But, first, this is precisely the answer we should expect from a theory of subjective reasons, second, this verdict seems to me the precisely correct answer, and, third, this does not shed any negative light on the general claim that my proposal, rather than the traditional view, *can*—though needn't always—say that it is subjectively wrong to proceed at speed even when one believes that no one is coming.

To conclude my alternative view, the subjective 'ought' is explained in light of EP1 and EP2. One subjective 'ought' to conform to the strongest subjective reasons, and these reasons are explained by objective reasons, in terms of which the genuine moral valence of our actions (i.e., the objective 'ought') is characterized.

5. *Objection: Epistemic Asymmetry*

An objection to my proposal brings up an issue on which I've briefly touched already. EP2 does not index a person's subjective reasons to her moral beliefs. In other words, subjective reasons are determined by $OW_{r,\phi}$, rather than $CRE(OW_{r,\phi})$.³² But this may seem an unmotivated asymmetry in the relevance of an individual's epistemic circumstances. Why should my subjective reasons be determined by my beliefs in the truth of an individual corresponding objective reason, but not my beliefs concerning the the moral significance of that corresponding objective reason?

I offer two thoughts in defense of this asymmetry. The first I leave until §6.3—I hold that accepting the epistemic asymmetry helps to avoid a violation of 'ought' implies 'can'. But I offer a more general defense here. To hold that one's subjective reasons are determined by one's beliefs about the world *along with* one's beliefs about the moral significance of what one believes would seem to tie one's subjective reasons to an agent's epistemic perspective in a way that generates absurd results. John Wilkes Booth, for instance, may have believed that the assassination of Abraham Lincoln would have succeeded in avenging the South, and that this is the case was a reason of overwhelming moral significance to assassinate Abraham Lincoln. If we are to hold that a person's subjective reasons are a function not only of the credence levels in corresponding objective reasons (i.e., "killing Abraham Lincoln will avenge the South for the loss of slaves"), but also their own assessment of the moral significance of these reasons (i.e., that killing Abraham Lincoln will avenge the South for the loss of slaves is a reason to kill

³²For a contrary view, see Zimmerman, 38-42.

Abraham Lincoln), we would be forced to claim that Booth had overwhelming subjective moral reason to kill Lincoln. This is implausible. (Notice that this implication holds only if the relevant epistemic circumstances are *belief*. But I think that merely changing the relevant epistemic circumstances to, say, *evidence*, or *idealized belief* is insufficient to avoid this problem. It seems plausible that Booth would have continued to believe that avenging the South was a moral reason to kill Lincoln even under a wide range of actual or idealized epistemic circumstances.) If—as OSM presumes—we wish to treat the subjective moral ‘ought’ as providing an (albeit imperfect) guide to the moral quality of an individual’s character, we surely cannot say that the mere fact that Booth accepted an objective moral reason to kill Lincoln helps to determine Booth’s subjective obligations. Furthermore, if conformity to one’s subjective moral obligations helps to determine the appropriateness of the reactive attitudes, or the appropriateness of blame, we would certainly not wish to index Booth’s subjective obligations to his moral beliefs. Doing so, in this case, would render him blameless for the death of Lincoln, surely a result to be avoided if at all possible.

Of course, it is natural in cases such as Booth’s to assume, e.g., that Booth believed that killing Lincoln might have accomplished some good end, such as promoting utility, etc.³³ We may be tempted to suggest that under this set of beliefs, Booth may have acted as he subjectively ought to have. But this is not the case we’re imagining: we are imagining that Booth believes that Lincoln’s murder will avenge for the loss of slaves is *itself* the reason to kill Lincoln. But if we assume that Booth has such mistaken beliefs, it is far less plausible to believe he behaved in a subjectively appropriate way. If so, the epistemic asymmetry is justified.³⁴

6. Three Objections to OSM in Light of EP1 and EP2

Answering the explanatory question with EP1 and EP2 rather than the traditional view is not only independently motivated, but pays substantial dividends for OSM. In this section, I discuss three traditional objections to OSM. I argue that my view, in contrast to the traditional view, can, if not avoid them altogether, significantly blunt their force. If so, I claim, this is good reason to accept my answer to the explanatory question, whether or not one goes on to accept OSM.

³³Thanks to [...].

³⁴See, for instance, Zimmerman, 12-14.

6.1. Information

OSM holds that subjective morality is “subsidiary” to objective morality. But Allan Gibbard disputes OSM’s priority rule:

It should be easy to see, indeed, that no definition of the subjective sense in terms of the objective sense will work. Suppose you are offered a bet on the flip of a coin: win a thousand dollars if the coin lands heads, and lose \$800 if it lands tails. You have no way of knowing how the coin will land. To judge what you ought to do objectively, we need to know how the coin will in fact land, but we don’t need to know much about the value of money: we just need to know that having more is better than having less. To decide what you ought to do subjectively, in contrast, we need to know more about value: how the gain of a thousand dollars compares to the loss of \$800. . . The onlooker, then, who knows objective *oughts* doesn’t thereby have the information he would need to settle what the subjective ought to do subjectively.³⁵

Gibbard’s argument runs like this. In deciding whether one is objectively required to take a chance on this coin flip, one needs to know only the actual outcome of the case, and that one outcome (i.e., winning \$1000) is better than another (losing \$800). But to know what one *subjectively* ought to do, it is not enough simply to know whether winning \$1000 is better than losing \$800. To know whether one objectively ought to flip the coin, one needn’t know that the loss of \$800, for instance, would render one destitute and that the gain of \$1000 would be enough, say, for a trip to fabulous Las Vegas. All one must know is that more is better. But that losing \$800 would render one destitute and that winning \$1000 would be enough for a trip to Vegas is surely of the essence in determining whether one subjectively ought to take a chance on this coin flip. Gibbard concludes that the subjective ‘ought’ must be prior, violating OSM.³⁶

If the explanatory question is answered by the traditional view, Gibbard’s analysis shows that OSM fails: the requisite objective moral requirements cannot explain subjective moral requirements; objective moral obligations offer too thin an explanation. But EP1 and EP2 can preserve the explanatory priority of objective morality—and hence OSM’s priority rule—in the face of Gibbard’s argument. To know whether I subjectively ought

³⁵Gibbard, 345.

³⁶Gibbard, 344.

to risk my money on this coin flip, I require information about the effects of winning and losing this coin flip. But, on EP2 (and as a matter of common sense), the information we require is information about *objective reasons*, and the importance of this information is given by its *objective* importance: to know whether I should gamble I must know the *objective* importance of the loss of \$800 (i.e., the objective significance of avoiding destitution) in comparison to the objective importance of the gain of \$1000 (i.e., the objective significance of going to Vegas). Insofar as to avoid destitution is far more objectively important than going to Vegas, one shouldn't, in this case, gamble (given the analysis offered by EP2). Here the explanatory priority of objective reasons is preserved, if not the objective 'ought'.

To put this another way, if my view is correct, to properly assess one's subjective 'ought', one must have a grasp of one's subjective reasons. But the subjective reasons for and against gambling are a function of propositions—such as “losing \$800 will cause destitution”, “gaining \$1000 will allow a trip to Las Vegas”, etc.—that do not contain an epistemic operator (i.e., *objective* reason-constitutive propositions). Thus one's subjective reasons are explained, at least in part, by corresponding *objective* reasons for and against such states. Hence though the objective 'ought', which doesn't require us to know much about these individual states, cannot explain the subjective 'ought', the information required to make subjective decisions is, according to EP2, a matter of objective *reasons*. But (a) this verdict is straightforwardly plausible, and (b) this preserves the extent to which objective and subjective morality obey OSM's priority rule: when we wish to determine the genuine moral valence of actions, we assess these actions *objectively*. None of this violates Gibbard's correct insistence that the objective 'ought' is not explanatorily prior to the subjective 'ought'.

6.2. Jackson Cases

Look again at the case of the miners. In the reasoning detailed above, (1) seems quite obviously correct: the best moral decision (i.e., the subjectively right act) in this case is to block neither shaft. But, if we accept that one *objectively* ought to block either A or B, the following problem arises. It seems platitudinous to claim that one does not make a morally appropriate decision if one behaves in a way one *knows* is morally wrong. But then if the objective 'ought' sets the genuine moral valence of actions, it would seem that blocking neither shaft is *not* the best decision in this case. After all, it is always morally inappropriate to perform an action one knows to be morally sub-optimal, and we *know* that blocking neither shaft is not, as a

matter of objective morality, morally sanctioned. But if we accept, as we certainly should, that the right decision in this case is to block neither shaft, we must reject OSM; the objective ‘ought’ does not determine the genuine moral valence of our actions. Otherwise a fully morally conscientious person could make a decision that he or she knows full well is morally sub-optimal. Call this the objection from—as they have come to be known—“Jackson cases”.³⁷

The argument from Jackson cases can be stated more or less formally as follows:

1. A morally conscientious person would not perform an action she knows to be morally wrong.
2. In the case of the miners (and other Jackson cases), the morally conscientious person would choose an action that she knows is not objectively morally permitted.
3. Hence, by (1) and (2), acting in a way that is not objectively permitted is not morally wrong.
4. Hence, whether an action conforms to the objective ‘ought’ does not determine the genuine moral valence of that action.

The independent plausibility of (1) and (2) has led some to reject the claim that the moral rightness of actions depends on the actual consequences of such actions. Indeed, Jackson rejects OSM on the basis of these cases, in favor of a view that identifies moral rightness with the action that maintains the greatest “expected moral utility”.³⁸

The proponent of OSM must deny (1) or (2). But insofar as (2) is surely correct, this leaves (1). And, indeed, the denial of (1) is itself an independently plausible strategy of response. As Krister Bykvist writes: “Since a moral agent should also be sensitive to the fact that in this particular case the value difference between a minor wrong-doing and a major one is greater than the value difference between a right-doing and a minor wrong-doing, his preference for a minor wrong-doing over a major wrong-doing should be *stronger* than his preference for a right-doing over a minor wrong-doing.”³⁹ But herein lies the problem. If we accept the traditional view, the denial of (1) is barred. After all, if adequate moral decisionmaking—as determined by subjective moral obligations—is tied very closely to our beliefs about actions

³⁷Cf. Jackson, 462-463. See also Zimmerman, 18; Holly M. Smith, “The ‘Prospective View’ of Obligation” in *Journal of Ethics and Social Philosophy* (2011).

³⁸Jackson, 464-7.

³⁹“How to Do Wrong Knowingly and Get Away With It” MS.

that are objectively required, then it seems to follow that moral conscientiousness cannot deviate from that which one believes is morally required. And, if so, good moral decisionmaking requires one to block either shaft A or B. Hence the proponent of the traditional view is committed to the denial of (2), i.e., that the morally conscientious person would block one of the shafts. But this is a paradigmatic case of poor moral decisionmaking.

However, if we accept EP1 and EP2, OSM can deny (1). To see this, consider the structure of subjective reasons in this case as determined by EP2. Given the case as stated, one has credence level .5 that either A or B contains the miners. That blocking neither shaft will kill only one miner has a very high credence value (.9, let's say). Assume also that the objective weight in favor of saving all miners is 1, of killing only one miner, insofar as it is almost as good as saving them all, .9, and of saving no miners is zero. Given this, we can analyze the reasons for and against the relevant alternatives in the following way ("I maintain credence .5 that the miners are in shaft A": *sr1*; "I maintain credence .5 that the miners are in shaft B": *sr2*; "I maintain credence .9 that blocking neither shaft will kill one person": *sr3*):

$$\begin{aligned}
 W_{sr1,BlockShaftA} &= (.5*1) = .5 \\
 W_{sr1,BlockShaftB} &= (.5*0) = 0 \\
 W_{sr2,BlockShaftB} &= (.5*1) = .5 \\
 W_{sr2,BlockShaftA} &= (.5*0) = 0 \\
 W_{sr3,BlockNeitherShaft} &= (.9*.9) = .81
 \end{aligned}$$

It would appear that on EP2, there is equal subjective reason to block shafts A and B. However, the subjective reason in favor of blocking neither shaft is weightier than the reason to block either A or B, and hence one subjectively ought to block neither shaft. Despite the fact that one knows this will not be morally optimal, it is supported by the weightiest subjective reasons, and hence is required of appropriate moral decisionmaking. Thus OSM can sensibly deny the claim, stated here by Michael Zimmerman, that "conscientiousness precludes deliberately doing what one believes to be overall morally wrong."⁴⁰ If EP1 is correct, conscientiousness requires a weighting of subjective reasons. And, if EP2 is correct, one's subjective reasons require the decision to block neither shaft. This is compatible with the claim that the objective 'ought' determines the genuine moral valence of actions; the morally conscientious person—the person oriented toward objective moral reasons—will act in a way she knows is morally prohibited.

⁴⁰Zimmerman, 18.

I do not claim here to have invented a new method of response to Jackson cases. Indeed, my method of response bears a striking similarity to the “maximize expected moral value” view suggested by Jackson. My point is, rather, that if we reject the traditional view, we can retain OSM’s identification of moral rightness with the objective ‘ought’ while continuing to hold that adequate moral decisionmaking requires the blocking of neither shaft. This, it seems to me, is a significant result for the defenders of OSM, and is good reason to reject the traditional view in favor of EP1 and EP2. Indeed, I refrain from taking a strong stand on whether one should reject (1), or should instead accept (4), and deny OSM. I am interested only in the best formulation of OSM. Insofar as OSM must reject (1), it is important to offer a response to the explanatory question that could vindicate this proposal. EP1 and EP2 do this. The traditional view does not.

6.3. ‘Ought’ implies ‘Can’

It is generally held that ‘ought’ implies ‘can’—if I have an obligation (subjective or objective) to ϕ , it should be the case that I *can* ϕ . However, Zimmerman argues that on the traditional view, the subjective ‘ought’ violates this principle. This is for a simple reason. I might believe that I have the power to cure cancer. If this is correct, one would think that I subjectively ought to cure cancer. It just so happens, however, that I am deluded. I have no such ability. Hence subjective morality implies that I ought to do something I cannot do, violating ‘ought’ implies ‘can’.⁴¹ It is certainly possible, no matter the “epistemic circumstances” to which one’s subjective obligations are relativized, that these epistemic circumstances might grant credence to a claim that one can do something that one cannot, in fact, do. And if this is correct, one’s subjective obligations may tell in favor of the performance of this action given the traditional view. If, for instance, all the evidence indicates that I can leap tall buildings in a single bound (rare, but certainly possible), I may very well be required to do so if doing so would, for instance, save someone from a disastrous plummet from a skyscraper. This is, after all, what one objectively ought to do in light of these epistemic circumstances. This is true even if we assume that one is not *actually* objectively required to perform the action given an inability to do so. After all, the traditional view treats one’s subjective obligations as what the objective obligations *would be* in light of one’s beliefs, under whatever relevant epistemic circumstances.

⁴¹Zimmerman, 14.

It should be noted that Zimmerman is really objecting only to an understanding of morality that permits of one, subjective, ‘ought’. If, instead, we accept OSM and allow two ‘ought’s, we might be tempted to respond to this objection by biting the bullet: the availability of the objective ‘ought’ might generate less pressure to require the subjective ‘ought’ to imply ‘can’. After all, it may still be the case that the objective ‘ought’ does so. But I regard this as a pretty feeble response. It seems wrong to say, even if we allow an objective ‘ought’, that to avoid blame, for instance, or to make an adequate moral decision, one should do (or attempt to do) what one cannot do.

However, if we accept EP1 and EP2 we can allow that the subjective ‘ought’ implies ‘can’. The basic idea is this. The traditional view treats our subjective obligations as the result of a certain counterfactual, viz., what our objective obligations *would be* were our (in this case non-moral) beliefs, given whatever relevant epistemic conditions, true. But insofar as our non-moral beliefs might include false beliefs about our own abilities, in certain cases our objective obligations from the perspective of this counterfactual hold that we ought to do things we *actually* cannot do. But according to EP1 and EP2, our subjective obligations do not depend on what our objective obligations would have been had our beliefs turned out true, but they depend on our subjective reasons, which are in turn a factor of the *actual* objective reasons and their *actual* weight.

The solution works like this. To accept ‘ought’ implies ‘can’ for the objective ‘ought’, we must accept that objective *reasons* imply ‘can’, as well: there can be no objective reason to perform an action I cannot perform. Unless we accept this, it is possible that the greatest balance of objective reasons may favor an action that couldn’t be performed, and hence the objective ‘ought’ would not imply ‘can’. But if objective reasons imply can, because subjective reasons are determined by objective reasons, subjective reasons must imply ‘can’ also. Take an example. Given that objective reasons imply ‘can’, that leaping this tall building in a single bound will save Lois Lane is no objective reason for me to do it: I cannot (though it may very well be a reason for Superman to do it). But now imagine that I believe I can leap such a building in a single bound. Do I have a subjective reason to do so? Answer: absolutely not! Because I cannot leap the building, *that* doing so will save Lois Lane is a reason of *no objective weight whatever* for me to do so, given that objective reasons imply ‘can’. If so, given EP2, the subjective weight of the reason for me to leap this tall

building must also be zero.⁴² Because objective reasons imply ‘can’, so do subjective reasons, and hence (given EP1) so does the subjective ‘ought’.

EP2, and not the traditional view, avoids this problem by treating one’s subjective reasons as given not by what one’s objective reasons *would be* in light of one’s epistemic circumstances, but rather by treating one’s subjective reasons as an outcome of the true objective weight of a particular reason and one’s credence *in* that particular reason. Insofar as the actual objective weight of a reason in favor of an action one cannot perform is zero, there will be no subjective reason to perform it. If we accept the view, expressed in EP2, that one’s subjective reasons are determined by one’s credence levels in corresponding objective reason-constitutive propositions, and by the *actual* moral weight of those reasons, one needn’t end up with a conception of subjective morality that violates this basic principle.

Finally, I add one concluding thought. An individual’s beliefs about their abilities will naturally affect their beliefs about the objective weight of particular reasons. If I believe I can leap a tall building in a single bound, I will believe that the objective weight of the reason “that leaping this tall building in a single bound will save Lois Lane” is a reason of non-zero weight to do so. That my view can solve this problem is in part a result of its embrace of the epistemic asymmetry discussed in §5: EP2 treats subjective obligations as a function of credence levels in particular objective reasons, not as a function of beliefs about the *objective weight* of these reasons. Hence, or so I claim, the ability to accommodate ‘ought’ implies ‘can’ lends further weight to the epistemic asymmetry on display in EP2.

7. Conclusion

OSM urges the acceptance of not one but two moral perspectives: one that is sensitive to an agent’s epistemic circumstances, another that is not so sensitive. According to OSM, subjective and objective moral ‘ought’s can and should play distinct, but important, roles in our moral life and practice.

However, OSM does not, by itself, answer the explanatory question. In this paper, I have argued that one way to answer the explanatory question is independently plausible and succeeds at blunting the force of three important objections to OSM. Insofar as OSM is reflected—as it is in Hume’s discussion of Hamden—in our moral experience, and insofar as OSM solves purported paradoxes about ‘ought’, this is very good news, indeed.

I should briefly mention what I have not done in this paper. First, I

⁴²See note 23.

have not attempted to provide a comprehensive argument for OSM. I take it on faith that the reader will agree that there is something rather plausible about it, and that it holds attractions for a variety of reasons. My goal here is simply to articulate the best conception of OSM, a conception that leaves it less vulnerable to objections. Second, I have not attempted to address all important arguments against OSM, including all arguments against the moral relevance of the objective perspective, subjective perspective, or the more general possibility of dual categories of moral assessment. Many important objections remain. I urge those who would confront such objections to do so with EP1 and EP2 in mind: in so doing, much skepticism about OSM can be avoided.