



OBJECTIVE MORALITY, SUBJECTIVE MORALITY AND THE EXPLANATORY QUESTION

BY DALE DORSEY

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 6, No. 3 | AUGUST 2012

URL: WWW.JESP.ORG

COPYRIGHT © DALE DORSEY 2012

Objective Morality, Subjective Morality and the Explanatory Question¹

Dale Dorsey

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.²

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 coconspirators.

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.³ 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 the construction of subjective moral assessment of action *given* an objective assessment. I argue that a proper answer to the

¹ I would like to thank David Brink, Krister Bykvist, Janice Dowell, Erin Frykholm, Clayton Littlejohn and a variety of anonymous reviewers for their helpful comments.

² David Hume (1983/1778) *The History of England, Vol. 5*, Indianapolis: Liberty Fund, p. 574. I want to thank Erin Frykholm 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.

³ See Henry Sidgwick (1981/1907) *The Methods of Ethics*, 7th ed., Indianapolis: Hackett Publishing Company, pp. 394-95.

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 the best answer to the explanatory question and to emphasize the importance of my preferred answer for wider debates concerning the overall tenability of dual moral perspectives. The plan of the paper runs as follows. In §1, I introduce the distinction between objective and subjective morality 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 “oughts”). Beginning in §6, I argue that my conception of this relationship can blunt the force of three 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 circumstances.
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.⁴

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 “oughts” 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.

⁴ Notice 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 way to understand OSM, it is not the only way and, as I argue in the next section, it is false.

Why posit an independent subjective morality if we assume that the fundamental moral facts are objective? It seems to me that objective and subjective morality play importantly distinct roles in our moral lives. Objective morality, insofar as it is fundamental, determines whether our actions in fact conform to that which we are morally required to do. If OSM is correct, when asking “Did I do the right thing?” we are asking “Did I conform to my *objective* moral obligations?” But moral assessment goes beyond an assessment of whether or not one did the right thing. Conformity to our subjective obligations, though this does not guarantee that we will do the right thing in all cases, says something important about our *character*. Performing the subjectively right act seems to indicate that we are disposed, or oriented, in morally appropriate ways. 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 morally justified. This reason is indirect in the sense that its importance is derived from the importance of conforming to one’s objective obligations. 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 (despite Hume’s initial judgment that Hambden’s action was a “blameable extreme,” his behavior in light of his subjective obligations entails that there are “great reasons of alleviation”).

The connection between subjective obligations and moral character or blame must, of course, be mediated by a more specific theory of moral character and the appropriateness of blame, which I shall not offer here. Furthermore, I do not want to state the connection between these concepts too strongly. It is likely that any connection between virtuous decisionmaking, praiseworthiness, etc., and conformity to one’s subjective obligations will be imperfect.⁵ However, though this connection is imperfect, there remains good reason to take it seriously in assessing an individual’s decisionmaking. Only this perspective offers any evidence concerning the quality of a person’s actions in light of his or her beliefs. Thus subjective morality seems important even if *mere* conformity to one’s subjective obligations is not – as it surely is not – a complete theory of virtuous character or the appropriateness of praise and/or blame.

Though I do not seek to argue in favor of OSM here, it is worth investigating one set of considerations that seems to me to tell in favor of

⁵ For instance, 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 luck. Similar reflections may apply to blameworthiness. Thanks to an anonymous reviewer.

OSM and its priority rule (these are the most salient to me, though there may be many more). Consider:

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, in turn, of the moral propriety of our moral decision making, etc.). Our untutored reaction is likely to be that, though we made the right decision in light of the evidence (i.e., our action is subjectively right), we actually did the wrong thing – the *fundamental* moral facts tell in favor of the action's wrongness. But 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. Both considerations support not just the existence of two "oughts," but also OSM's priority rule. The reason that 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. Furthermore, in declaring that we did the *wrong thing* in punishing Dave, this appears to indicate that the genuine moral valence of our actions – the answer to the question "Did I do the right thing?" – is determined by considering our objective moral obligations. If this is correct, it would appear that the fundamental moral facts are objective; the subjective "ought" is thus subsidiary (though, perhaps, no less significant for our moral lives). Though this is certainly no knock-down argument in favor of OSM, I hope it shows that OSM is at least plausible *enough* to consider further.

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* (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, whichever 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.”

2. The Traditional View

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 the contents of one’s beliefs are true, one objectively ought to ϕ . (I use “beliefs” here simply for shorthand. Those who prefer subjective “oughts” to be relativized to a different epistemic state – such as evidence, knowledge, evidence that has been idealized in some way, etc. – are free to make the necessary translations. Nothing in the argument of this paper will hinge on the particular preferred epistemic state to which subjective obligations are indexed.)

This conception is expressed 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 ”;⁶ Zimmerman: “An agent [subjectively] ought to perform an act if and only if he believes that it is the best option that he has”;⁷ Derek Parfit, who writes (concerning the practical, nonmoral “ought”): “[W]e can call some act what we *ought practically* to do in the *evidence-relative* sense just when this act would be what we had decisive reasons to do, if we believe what the available evidence gives us decisive reasons to believe, and these beliefs were true”;⁸ and Peter Railton, who writes, describing “subjective

⁶ Niko Kolodny and John MacFarlane (2010) “Ifs and Oughts,” *Journal of Philosophy* 107: 118. Note that, for Kolodny and MacFarlane, one’s “best choice” (not relativized to one’s epistemic states) is one’s objective obligation. See Kolodny and MacFarlane: 117.

⁷ Michael Zimmerman (2008) *Living With Uncertainty*, Cambridge: Cambridge University Press, p. 5. Note that, for Zimmerman, one’s “best choice” is equivalent to one’s objective obligations. See Zimmerman: 2.

⁸ Derek Parfit (2011) *On What Matters, Vol. 1*, Oxford: Oxford University Press, p. 163. Note that, for Parfit, “ought” is a product of decisive reasons (33). Parfit also offers a “belief-relative” sense of the practical “ought.” For my purposes, both accounts display the “traditional” account of a subjective “ought,” i.e., that we subjectively ought to perform whatever action that, were the relevant epistemic state (such as evidence, beliefs) accurate, we objectively ought to perform.

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.”⁹ On Railton’s view, one subjectively ought to do that which in one’s best estimation will actually produce the best consequences, and hence what, in one’s best estimation, will be the subject of an objective “ought.” Though these views are obviously not equivalent, they all accept what I identify as the traditional view: that one’s subjective obligations are determined by what one’s objective obligations (or, what is for these thinkers the same: “best choice,” “best option,” “decisive reasons”) *would be* under the assumption that the contents of one’s relevant epistemic circumstances (which, as I note above, I will abbreviate simply to “beliefs”) are true. To put this another way, the traditional view holds that the explanatory priority of objective morality is at the level of moral *requirements*; subjective moral requirements are determined by whatever one’s objective requirements *would be* given one’s beliefs.

My statement of the traditional view requires one bit of disambiguation. Some versions of the traditional view, such as Zimmerman’s, will allow that the subjective “ought” is relative not just to a person’s nonmoral beliefs, but also relative to a person’s *moral* beliefs, i.e., beliefs concerning the moral facts.¹⁰ Kolodny and MacFarlane and Parfit seem to hold that a person’s subjective “ought” is sensitive only to an individual’s nonmoral beliefs. 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 she objectively ought to do on the assumption that her *nonmoral* beliefs are true.

This ambiguity notwithstanding, the traditional view is popular. Its popularity is certainly understandable. After all, if the subjective moral perspective informs evaluation of decision making or moral character, and if our decision making 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. The primary problem with the traditional view is that its interpretation of the priority of objective morality to subjective morality is cast at the level of the

⁹ Peter Railton (2003) “Alienation, Consequentialism, and the Demands of Morality,” in *Facts, Values, and Norms: Essays Toward a Morality of Consequence*, Cambridge: Cambridge University Press, p. 165.

¹⁰ For Zimmerman, the subjective “ought” is similar to Derek Parfit’s “normative-belief-relative” “ought”: “[W]e *ought practically* to do in the *normative-belief-relative* sense just what we believe that we ought practically to do, or what we believe that we have decisive reasons to do” (Parfit: 163).

moral “ought.” But this cannot be correct. 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 (I will say more about what these are in the following sections). The argument for EP1 runs as follows. 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 is 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,¹¹ or any other form of moral rationality one might choose to accept.¹²)

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 the 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

¹¹ See, for instance, Michael Slote (1989) *Beyond Optimizing*, Cambridge, Mass.: Harvard University Press.

¹² For a nonstandard account of moral rationality, see Douglas Portmore (2008) “Are Moral Reasons Morally Overriding?” *Ethical Theory and Moral Practice* 11(4): 369-88.

there is the greatest balance of objective moral reasons, but that it is 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. Consider a case similar to one offered by Allan Gibbard. Imagine that, in driving to a meeting, I approach an intersection at which there is a pedestrian crosswalk. Assume that I am 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 am not really very sure, I *believe* that there is no pedestrian coming. This case seems to me to show two things. First, it seems right to say that one is subjectively required to slow down. Gibbard shares this view: “At a blind intersection ... speeding on through is subjectively wrong, even if most likely no cross-traffic is coming.”¹³ However, it seems plausible to say that this does not exhaust what subjective morality has to say about this case. After all, there appears to be at least *some* subjective consideration in favor of proceeding at speed; after all, I do not believe that anyone is coming, and slowing down will inconvenience many. But if this is right, it is hard to see how subjective reasons could fail to exist: Even in the face of a subjective obligation to the contrary, one has subjective reason to proceed at speed.

So what does all this say about the traditional view? Offhand, it seems right to say that *if* one’s subjective obligations are explained in terms of subjective reasons, then it cannot be the case that, e.g., one’s subjective obligations are to perform that which would be the “best choice” in light of one’s beliefs. Instead, it seems right to say that one’s subjective obligations are to perform that which is supported by decisive subjective reasons. The objective “ought” (“best choice,” “best option” or “decisive reason”) does not appear to play much of a role. One might salvage the traditional view and hold instead that subjective moral *reasons* are themselves explained in terms of objective moral requirements.¹⁴ One might say, for instance, that one has subjective reason to ϕ if ϕ -ing would be objectively required given one’s epistemic circumstances. If this is correct, the traditional view follows: If one has subjective reason only to perform that action that is objectively required given one’s epistemic circumstances, it follows trivially that one subjectively ought to perform that action that would be objectively obligated under such

¹³ Allan Gibbard (2005) “Truth and Correct Belief,” *Philosophical Issues* 15: 345.

¹⁴ This 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” (2003) “Prudence and Authenticity: Intrapersonal Conflicts of Value,” *The Philosophical Review* 112: 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.

conditions. But Gibbard's example shows that this suggestion also fails. If this were correct, I would only have subjective reason to proceed at speed, insofar as I *believe* no one is coming. In light of this belief, I objectively ought to proceed at speed: The "best choice," "best option," or that which I have "decisive reason to do," in light of my belief, is not to stop. But if this is the only subjective reason I have, EP1 would hold that I am subjectively required to proceed at speed. But I am not. And hence it would appear that I can have subjective reason – even decisive subjective reason – to act in a way contrary to that which would be objectively required in light of my beliefs or wider epistemic circumstances. This conclusion sinks the traditional view. It seems right to say, in many cases, that I am subjectively required to do – I have decisive subjective reason to do – something *other* than that which I would be objectively required to perform *were* that which I believed to come out true. But the traditional view denies this, and hence the traditional view is false.

A response to EP1 should be considered here.¹⁵ 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.¹⁶ 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 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 would not wish to deny that the fact that ψ -ing has positive (though nonoptimal) 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 (subjective) reason to ψ ; but this reason could be outweighed by weightier (subjective) reasons – understood in terms of expected moral value – in favor of ϕ -ing.¹⁷

¹⁵ Thanks to an anonymous reviewer for suggesting that I consider this view.

¹⁶ Cf. J. J. C. Smart (1974) "An Outline of a System of Utilitarian Ethics," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press, ch. 7.

¹⁷ Note 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*.

3. An Alternative View: Part 1

To reject the traditional view may seem to put OSM in jeopardy. After all, if the objective “ought” is not explanatorily prior to the subjective “ought,” this would appear to void OSM’s priority rule. OSM, after all, declares that objective morality is explanatorily prior. But there is more to morality than obligations (or “best choices,” “best options” or “decisive reasons”). 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 explained 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 counts in favor of 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 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. He 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 .¹⁸

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 do not 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 do not believe that there is a pedestrian coming, I still seem to have a reason to slow down *given* that I cannot 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 could not have subjective reason to slow down, on Schroeder’s view.¹⁹

¹⁸ Mark Schroeder (2007) *Slaves of the Passions*, Oxford: Oxford University Press, pp. 13-14.

¹⁹ An anonymous reader indicates that this conclusion might not hold if (a) one believes that a pedestrian *might* be coming and (b) if “a pedestrian might be coming” would, if true, be

I want to tweak Schroeder's account along two dimensions. First, it seems to me more plausible to treat subjective reasons not as the believed proposition, but rather the proposition that *expresses* the agent's relevant epistemic state.²⁰ Second, it seems sensible not to limit the epistemic state to "belief," insofar as we have already seen that disbelieved propositions might themselves constitute subjective reasons to slow down. 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.²¹

*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 maintains some epistemic valence for an agent and *would* count as an objective reason in favor of a particular action *were it true*.²² "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"

objective reason to slow down. However, this seems to me an odd way of describing what is going on in such cases. First, it sounds odd to say that a claim about probability is an objective reason, insofar as objective reasons are determined by what is actually the case, not what *might* be the case. Leaving this aside, however, to believe that someone might be coming just is to grant a certain credence level to the proposition "someone is coming." But if this is right, the subjective reason itself is better described by *Subjective**: a subjective reason is a corresponding objective reason ("someone is coming") placed within an epistemic operator, in particular, a credence level.

²⁰ See Mark Schroeder (2008) "Having Reasons," *Philosophical Studies* 139: 60-61. One reason to accept this view is that the alternative implies that that *false* propositions can be reasons (i.e., subjective reasons). But this sounds strange to me: How could a proposition be a reason of any kind if it is false? Schroeder himself accepts this conclusion (2008: 14, n.23), so I refrain from treating this as an objection to his view here. In any event, 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.

²¹ I 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 were p true, p would constitute an objective reason for x to ϕ .

²² Thus *Subjective** is not a "Factoring Account" of subjective reasons. In other words, my view is not that "the subjective reason relation is merely a restriction on the objective relation" (Schroeder (2008): 60). Corresponding objective reasons can be false, and hence can fail to be genuine objective reasons. But subjective reasons are based on the objective weight of these reasons on the assumption that these reasons are, in fact, true.

has a positive epistemic status for me, which renders it the case that I have a subjective reason to punish Dave.²³

4. An Alternative View: Part 2

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 corresponding objective reason) will maintain *some* epistemic status or other for me, 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, assuming that "Dave is guilty" is true. 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. (Perhaps after due reflection, see note 30.) 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 is of

²³ "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.

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 p for a given individual A is determined by the strength of A 's belief in p .²⁴ A credence level of 1 indicates that A maintains the strongest possible belief in p . A credence level of 0 means that A grants p no credence whatever: A maintains the strongest possible *disbelief* of p . A credence level of .5 indicates that A is epistemically indifferent to p .²⁵

In light of this terminology, it seems sensible to determine the weight of a subjective reason as a function of the objective weight of the corresponding objective reason, and of my credence level *in* this corresponding objective reason. 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 .²⁶ A natural possibility, however, is multiplication.²⁷ Take an easy case. If we assume that the objective weight of "Dave is guilty" 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

²⁴ Those 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.

²⁵ One might put this in terms of one's assessment of the objective chance of a proposition: A credence level of 0 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 the proposition an equal chance of being true or not true. A credence level of 1 means that you grant the proposition no chance of being false. See, for instance, David Lewis (2000) "A Subjectivist's Guide to Objective Chance," in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press.

²⁶ One desideratum seems sensible. If it is 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 0, despite its heavy objective weight, this should also provide no subjective reason.

²⁷ One 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.

is guilty” ($sr1$): $SW_{sr1, PunishDave} = (.51 * 1) = .51$.²⁸ Compare this to the weight of “I doubt Dave is innocent” ($sr2$) as subjective reason not to punish Dave. Assuming that one grants “Dave is innocent” credence value .49, and that the objective weight of “Dave is innocent” in providing reason not to punish Dave is 1: $W_{sr2, Don't PunishDave} = (.49 * 1) = .49$. 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, I 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 the relative moral importance of inconvenience). Assume that the respective credence values are .3 and .7, respectively. If so, the weight of “I doubt someone is coming” ($sr1$) and “I believe there is no one 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.²⁹

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. If that is right, there may be subjective reason enough to proceed at speed. But this point is neither here nor there (and, in fact, should be precisely what we would expect from a proper theory of subjective reasons). After all, the problem with the traditional view is that it cannot seem to say under any circumstances that one ought to slow down if one believes, no matter how weakly, that no one is coming. (This is also an upshot of Schroeder’s theory of subjective reasons: Only *believed*, rather than *disbelieved*, propositions can count as subjective reasons.) But my view avoids this problem. Depending on one’s credence level, and the relative moral significance of preventing

²⁸ By 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.

²⁹ Indeed, 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 25) 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.

inconvenience and preventing serious injury, EP1 and EP2 correctly hold that one can be subjectively required to slow down *even if* one believes rather strongly that no one is coming. And it can maintain this verdict while treating objective morality – in particular, the objective weight of corresponding objective reasons – as explanatorily prior to subjective reasons and hence the subjective “ought.”³⁰

To conclude my alternative view, the subjective “ought” is explained in light of EP1 and EP2. One subjectively “ought” to conform to the strongest subjective reasons, and these reasons are explained by the objective weight of corresponding objective reasons, according to the formula in EP2.

5. Objection: Epistemic Asymmetry

An objection to my proposal brings up an issue on which I have briefly touched already. EP2 does not index a person’s subjective reasons to her moral beliefs. In other words, subjective reasons are determined by $OW_{\neg\phi}$, rather than $CRE(OW_{\neg\phi})$.³¹ 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 concerning the truth of a corresponding objective reason, but not my beliefs concerning the *objective 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 was a reason of overwhelming moral significance to assassinate Abraham Lincoln. If we are

³⁰ An anonymous reviewer complains on the following grounds. Imagine that I have never considered the thesis, “My standing on top of the Empire State Building will cure cancer.” If true, however, this certainly provides a very strong reason to stand on top of the Empire State Building. And given that I have never considered it, it would seem that the proper credence is indifference: .5. If so, it would seem that I have a comparatively very strong subjective reason to stand on top of the Empire State Building given the moral significance of a cure for cancer. But this seems wrong. One reasonable proposal in response seems to me to treat an individual’s credence levels not as simply whatever their belief-states happen to be at a particular time, but rather what their credence level would be after due reflection. If, after due reflection, I am genuinely indifferent about whether standing on top of the Empire State Building will cure cancer, it seems to me precisely the correct verdict to say that I have a reasonably strong *subjective* reason to do so. I do not believe it will, but I do not believe it *will not*, either. Most, however, will treat this hypothesis with a strong degree of suspicion. Given this, though there will be *some* reason to stand on top of the Empire State Building, it will be comparatively weak, indeed.

³¹ For a contrary view, see Zimmerman: 38-42.

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"), but also their own assessment of the moral significance of these reasons (i.e., that killing Abraham Lincoln will avenge the South is a reason to kill 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. It seems right to say that part of what renders John Wilkes Booth so blameworthy *is* his mistaken moral beliefs. Had he simply possessed mistaken factual beliefs (i.e., that killing Lincoln would have prevented the rise of Jim Crow laws), we would be less likely to blame him; we would, as Hume does when considering Hamden, allow "great reasons of alleviation."

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. 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 are imagining: We are imagining that Booth believes that "Lincoln's murder will avenge the South" 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.³²

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

³² See, for instance, Zimmerman: 12-14.

good reason to accept my answer to the explanatory question as the best statement of OSM, whether or not one accepts OSM, all things considered.

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 subject ought to do subjectively.³³

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 \$1,000) is better than another (losing \$800). But to know what one *subjectively* ought to do, it is not enough simply to know whether winning \$1,000 is better than losing \$800. To know whether one objectively ought to flip the coin, one need not know that the loss of \$800, for instance, would render one destitute and that the gain of \$1,000 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 \$1,000 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 goes on to conclude that the subjective “ought” must be prior, violating OSM.³⁴

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 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 the objective importance of certain facts or propositions: To know whether I should gamble I must know the *objective* importance of the loss of \$800 (i.e., the objective significance of

³³ Gibbard: 345.

³⁴ Gibbard: 344.

avoiding destitution in providing a reason not to gamble) in comparison to the objective importance of the gain of \$1,000 (i.e., the objective significance of going to Vegas as a reason to gamble). Insofar as to avoid destitution is far more objectively important than going to Vegas, one should not, in this case, gamble (given the analysis offered by EP2). Here we preserve the explanatory priority of objective morality – though not the objective “ought” – while suffering no paucity of information.

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 \$1,000 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 does not 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 this verdict (a) is straightforwardly plausible and (b) preserves the extent to which objective and subjective morality obey OSM’s priority rule: When we wish to determine the answer to the question “Did I do the right thing?” we assess our 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

Consider 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.³⁵

In the case 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 equally obvious is that one objectively ought to block either A or B; blocking neither shaft will guarantee a morally suboptimal outcome. But if we accept that one *objectively* ought to block either A or B, the following problem arises for OSM. 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

³⁵ Kolodny and MacFarlane: 115.

then if the objective “ought” sets the fundamental 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 is not the right thing, and we *know* that blocking neither shaft is not, as a matter of objective morality, the right thing. 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 whether, in fact, we did the right thing. Otherwise a fully morally conscientious person could make a decision that he or she knows full well is *not* the right thing. Call this the objection from – as they have come to be known – “Jackson cases.”³⁶

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; it does not determine whether the agent “did the right thing.”
5. Hence, objective morality is not explanatorily prior.³⁷

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

³⁶ Cf. Frank Jackson (1991) “Decision-Theoretic Consequentialism and the Nearest and Dearest Objection,” *Ethics* 101: 462-63. See also Zimmerman: 18; Parfit: 159-60; Holly M. Smith (2011) “The ‘Prospective View’ of Obligation,” *Journal of Ethics and Social Philosophy*.

³⁷ (5) follows from (4) on an important assumption about OSM: that OSM’s priority rule is motivated by the thought, as explored above, that the fundamental moral facts are objective rather than subjective. In any event, I will grant, for the sake of argument, that (5) follows from (4).

³⁸ Jackson: 464-67.

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 moral conscientiousness – as determined by subjective moral obligations – is tied very closely to our beliefs about actions that are objectively required, then it seems to follow that moral conscientiousness cannot deviate from that which one believes is morally required.⁴⁰ And, if so, good moral decision making 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, or so it would seem, is precisely what the morally conscientious person would *not* do.

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, 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, is .9; and of saving no miners is 0. 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 subjectively required. 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.”⁴¹ If EP1 is correct,

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

⁴⁰ Alternatively, the partisan of the traditional view could hold that moral conscientiousness is possible by doing that which one knows is objectively wrong *and* failing to conform to one’s *subjective* obligations. But this position is surely untenable.

⁴¹ Zimmerman: 18.

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.

I do not claim here to have invented a new method of response to Jackson cases. Indeed, my method 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 decision making 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."⁴² 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

⁴² Zimmerman: 14.

objective obligations *would be* in light of one's beliefs, under whatever relevant epistemic circumstances.⁴³

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 nonmoral) beliefs, given whatever relevant epistemic conditions, true. But insofar as our nonmoral 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 to be true, but they depend on our subjective reasons, which are in turn a factor of the genuine objective weight of corresponding objective reasons.

To see why this matters, consider the following. To accept "ought" implies "can" for the objective "ought," we must accept that objective *reasons* imply "can," as well: No proposition can count as a reason for A to ϕ unless A can ϕ . Unless we accept this, it is possible that the greatest balance of objective reasons may favor an action that could not be performed, and hence the objective "ought" would not imply "can." But if objective reasons imply can, because subjective reasons are determined by the objective weight of corresponding objective reasons, subjective reasons must imply "can" also. Take an example. Say that I come to believe that leaping this tall building in a single bound will save Lois Lane from certain death. Is "I believe that leaping this tall building in a single bound will save Lois Lane from certain death" a subjective reason to leap the tall building? To determine the answer to this question, we must know the objective weight of "leaping this tall building in a single bound will save Lois Lane from certain death" as a reason to leap the tall building in a single bound. But, quite clearly, even if this proposition is true, it provides no objective reason for me to do so; the objective weight of this proposition must be 0. Why? Because objective reasons imply "can," and I cannot. 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? Again, absolutely not. Because I cannot leap the building, *that* doing so will save Lois Lane is a corresponding reason of *no objective weight whatever* for me to do so, no matter what I believe about my own abilities. If so, given EP2, the subjective weight

⁴³ 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 "oughts," 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.

of the reason for me to leap this tall building must also be 0.⁴⁴ Because objective reasons imply “can,” so do subjective reasons, and hence (given EP1), so does the subjective “ought.”⁴⁵

Finally, I add one concluding thought. People’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 EP2.⁴⁶

⁴⁴ See note 26.

⁴⁵ An anonymous reviewer complains. What about the proposition: “I can leap this tall building in a single bound and leaping this tall building in a single bound will save Lois Lane from certain death.” If true, would this entail that I have an objective reason to leap the tall building? Of course! Now consider, “I believe that I can leap this tall building in a single bound and leaping this tall building in a single bound will save Lois Lane from certain death.” Is this a subjective reason to do so? Given the fact that, if true, the former would be an objective reason, this had better be a subjective reason (given my strong credence level). And so, even if I *cannot* leap the tall building, I have subjective reason to do so. But this analysis is mistaken. This analysis treats the entire believed conjunction *as* the corresponding objective reason. But that is not right. The right-hand conjunct, in this case, is simply a necessary condition for there to be any reason at all. It *itself* is not a reason. And so, given EP2, the weight of the corresponding objective reason in favor of leaping the tall building is 0, given that the reason (i.e., the left-hand conjunct) does not apply (given that I *cannot* leap the tall building). And though “I can leap this tall building in a single bound and leaping this tall building in a single bound will save Lois Lane from certain death” would certainly imply an objective reason to do so if true, this does not entail that this proposition itself *constitutes* a corresponding objective reason (which would then imply a subjective reason were I to believe it). Rather, it constitutes a reason (left-hand conjunct) and a necessary condition for the reason (right-hand conjunct).

⁴⁶ An anonymous reviewer suggests that to insist that subjective obligations imply “can” may get in the way of the connection between blameworthiness and conformity to one’s subjective obligations. Imagine, for instance, that I believe that I can save this drowning child. As it happens, however, I cannot. As it turns out, however, I just stand there, rather than making an effort despite my beliefs. Is it plausible to say that I am blameless in not even making an effort, despite the fact that I had no subjective moral obligation to save the child? Can it be plausibly said that I am properly morally oriented? I think the answer to these questions is no, but this fact does not mean that there is no connection between blameworthiness and subjective obligations. It seems quite right to say that, though I have no subjective obligation to save the child, I may very well have a strong subjective obligation to *try* to save the child. Here is why I think this: “I believe I can save this drowning child,” while perhaps not constituting a subjective reason to save the child in a situation in which I cannot, may nevertheless constitute an *objective* moral reason to *try* to save the child. And hence, it seems right to say that I will have a strong subjective moral reason to try to do so (given that my credence in that proposition will, very likely, approach 1). In a case in which I do not try to save the child, then, I am blameworthy. Not because I did not conform to my subjective

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 "oughts" 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 Hamlden – in our moral experience, this is very good news indeed.

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

Dale Dorsey
University of Kansas
Department of Philosophy
ddorsey@ku.edu

obligation to save the child, but because I did not conform to my subjective obligation to *try* to do so. Of course, as noted before, all of this depends on substantive theories of blameworthiness and of the nature of one's moral obligations. But, for what it is worth, this seems to me a plausible analysis of the considered judgment that I am blameworthy in this case, and maintains the relevant (albeit imperfect) connection between subjective morality and blameworthiness.