[Forthcoming in *The Journal of Ethics*]

A Defense of the Objective/Subjective Moral Ought Distinction

Kristian Olsen

Abstract: In this paper, I motivate and defend the distinction between an objective and a subjective moral sense of "ought." I begin by looking at the standard way the distinction is motivated, namely by appealing to relatively simple cases where an agent does something she thinks is best, but her action has a tragic outcome. I argue that these cases fail to do the job—the intuitions they elicit can be explained without having to distinguish between different senses of "ought." However, these cases are on the right track—I argue that more sophisticated versions of the cases provide strong motivation for the distinction. I then discuss two important problems for the distinction: the "which 'ought' is more important?" problem, and the "annoying profusion of 'oughts'" problem. I argue that each of these problems can be solved in several different ways.

1. Introduction

The term "ought" is ambiguous. That much is uncontroversial. If you say that I ought to tip 20%, this could mean that I *morally* ought to tip 20%, that I *legally* ought to tip 20%, that I *prudentially* ought to tip 20%, that I *etiquettically* ought to tip 20%, and so on.

Even if we limit our attention to the moral sense of "ought," the term is still uncontroversially ambiguous. If you say that I morally ought to keep my promise, this might mean that I *pro tanto* morally ought to keep my promise, or that I *overall* morally ought to keep my promise.[1] To say that I pro tanto morally ought to keep my promise is to say, roughly, that there is a moral reason for me to keep my promise. To say that I overall morally ought to keep my promise is to say, roughly, that the moral reasons in favor of keeping my promise outweigh the moral reasons against it.

If we limit our attention to the overall moral sense of "ought," do any ambiguities remain? Many philosophers say "yes": we need to distinguish between an *objective* and a *subjective* sense of the term.[2] According to these philosophers, to

[1] This distinction is stressed by W. D. Ross (1930; 1939).

[2] See, e.g., Ewing (1948: 112-47), Brandt (1959: 360-67), Russell (1966), F. Feldman (1986; 2012), Portmore (2011), and Dorsey (2012). I cite several additional examples in the next section.

say that I overall objectively morally ought to do something is to say, roughly, that doing it is overall most morally appropriate in light of the facts of my situation. On the other hand, to say that I overall subjectively morally ought to do something is to say, roughly, that doing it is overall most morally appropriate in light of my beliefs about (or, perhaps, my evidence concerning) the facts of my situation.

Why do many philosophers think we should distinguish between an objective and a subjective sense of "overall morally ought"?[3] Typically, the distinction is motivated by appealing to relatively simple hypothetical examples where an agent does something she believes (or has good reason to believe) is best, but her action in fact has a terrible outcome. In section 2, I discuss several cases of this sort and argue that they do not sufficiently motivate the distinction between an objective and a subjective sense of "ought"; the cases can be accommodated by appealing to the distinction between obligation and other normative concepts that are universally accepted—in particular, blameworthiness and value. However, I do think the distinction ultimately should be drawn, and in sections 3 and 4, I develop more sophisticated cases that provide stronger motivation for it. Finally, in section 5, I defend the distinction from two of its most important problems: the "annoying profusion of 'oughts'" problem, and the "which 'ought' is more important?" problem.

2. The Typical Case for Drawing the Distinction

Those who accept the distinction between an objective and a subjective sense of "ought" typically motivate the distinction by appealing to a hypothetical example like the one Frank Jackson describes in an early paper on the topic. Jackson asks us to suppose that his daughter has tonsillitis and that he agrees to allow her to have a tonsillectomy. Tragically, however, the child dies during the surgery. Did Jackson do the right thing or the wrong thing? He answers,

> Obviously, the correct response is not to settle on one or the other reply but to distinguish. I did the subjectively right thing. I desired good and chose a means very likely to achieve it. I would have failed in my duty had I done otherwise. Nevertheless bad, not good resulted. I did the subjectively right but objectively wrong thing. (1986: 351)

Graham Oddie and Peter Menzies say similar things, but they use the example of Oedipus to make their point. Oedipus killed his father and married his mother, but he did not realize it—in fact, he was specifically attempting to avoid doing these things. Oddie and Menzies claim that

[3] From now on, unless otherwise noted, when I use the term "ought," I will always have the overall moral sense in mind.

> Cases like this lead us naturally to a familiar distinction: between objective and subjective notions of moral rightness. (1992: 512)

Likewise, Holly Smith considers the case of Tom—a hypothetical security guard in the Twin Towers during 9/11—who directs office workers to evacuate the building via the stairs instead of the elevator. Sadly, the stairs take too long and the workers die when the building collapses; however, if they had taken the elevator instead, they would have escaped. Smith remarks:

> When we focus on the actual outcome of Tom's action of telling the employees to use the stairwell, we want to say that it is wrong—indeed tragic, since it results in the avoidable deaths of scores of office workers. But if we focus on what Tom reasonably believes about the employees' options at the time he advises them, we want to say that he did the right thing…To resolve the apparent paradox of such actions being judged both right and wrong, moral theorists in the first half of the twentieth century argued that we must recognize several different senses of moral terms such as "right," "wrong," "obligatory," and "permissible." In one sense of "morally right," they argued, we mean something like "the morally best action in the actual circumstances," where in another sense of "right," we mean something like "the action that is morally most appropriate in light of the agent's beliefs about those circumstances, even if those beliefs are mistaken." (2010: 65)

Many other philosophers use similar examples to motivate the distinction between an objective and a subjective sense of "ought"/"right"/"wrong."[4]

In light of these passages, the standard way of motivating the distinction between an objective and a subjective sense of "ought" seems to be this. A case is given where an agent does something she thinks (or, at least, has good reason to think) is best, but the action has unexpected, terrible consequences. The case is supposed to evoke in us a tension: part of us wants to say that the agent has done what she ought to do, but part of us also wants to say that the agent should have done something else. And, the thought goes, the best way of accommodating this tension is to distinguish between an objective and a subjective sense of "ought." If we make such a distinction, then we can say that it's *true* that the agent did what she ought to do (when we have the subjective sense in mind), but it's *also* true that the agent didn't do what she ought to have done (when we have the objective sense in mind).

While many philosophers distinguish between an objective and a subjective sense of "ought," not all do. Some—the Objectivists—think that "ought" is univocal (in its overall moral sense) and objective; that is, they think it is determined

[4] See, e.g., Pollock (1979: 109-10), Gibbard (2005: 340), and Parfit (2011: 151-52).

by the facts of an agent's situation. Others—the Subjectivists—think that "ought" is univocal (in its overall moral sense) and subjective; that is, they think it is determined by the agent's beliefs about (or, perhaps, evidence concerning) the facts of her situation.[5]

Neither Objectivists nor Subjectivists should be impressed by the standard argument for "Sense-Splitting," as we might call the view. Let's focus on Holly Smith's Twin Towers case. In response, Objectivists can deny that what seems intuitive is that Tom's action is both right and wrong. Instead, what is intuitive is that both positive and negative things should be said about his action, morally speaking. But this, Objectivists can maintain, is perfectly compatible with their view, for they can claim—negatively—that when Tom tells the office workers to take the stairs, he acts wrongly (since this in fact leads to their death); however—positively—he is *blameless* for his action (since he has little reason to believe that it will have such bad results). Objectivists can therefore insist that they can provide both a positive and a negative moral evaluation of Tom's action and thereby accommodate the intuitive tension it evokes.

Subjectivists can respond to the case in a similar way. They, too, can claim that what is intuitive about the case is not that Tom acts both rightly and wrongly but that Tom's behavior should be evaluated in both a positive and a negative moral way. But like Objectivists, Subjectivists can maintain that they have the resources to say both positive and negative things about Tom's action. They can hold—positively—that when Tom tells the office workers to take the stairs, he does what he ought to do (since his evidence suggests that this is the safest escape route). However, they can also acknowledge—negatively—that Tom's action has terrible consequences; other options would have had drastically better results. In this way, Subjectivists can argue that they can account for the intuitive tension the case generates.

If successful, the Objectivist and the Subjectivist responses to the Twin Towers case look superior to the Sense-Splitter response. The moral landscape is conceptually rich—there are many ways of morally assessing agents and actions. It is undeniable that deontic concepts (obligatoriness/rightness/wrongness), axiological concepts (goodness/badness), and responsibility concepts (praiseworthiness/blameworthiness) populate the moral landscape. Even Sense-Splitters will agree with this. But the Objectivist and the Subjectivist responses contend that the moral landscape is rich enough *as it already is* to accommodate the intuitive data

[5] Objectivists include Moore ([1903] 1993; [1912] 2005), W. D. Ross (1930), Thomson (1986), Zimmerman (1996), and Graham (2010). Subjectivists include Prichard ([1932] 2002), W. D. Ross (1939), Howard-Snyder (2005), Zimmerman (2006; 2008), and Mason (2013). Another option: one might hold that "ought" is univocal (in its overall moral sense) and context sensitive (see Kolodny and MacFarlane [2010]). I will set this view aside in this paper.

elicited by the Twin Towers case (and other cases like it). There is no need, these approaches claim, to further enrich the moral landscape by drawing a distinction between an objective and a subjective ought.[6] If they are correct about this, then parsimony considerations render these approaches superior to the Sense-Splitter response.

However, while the Objectivist and the Subjectivist responses to Twin Towers look attractive at first glance, they are ultimately inadequate. In the next section, I will take a closer look at the Subjectivist approach. In the section after that, I will discuss the Objectivist approach. We'll see in these sections that Objectivism and Subjectivism share a similar problem. For each view, I will present a hypothetical example that looks difficult to accommodate in the usual way (that is, by appealing to the distinction between obligation and blameworthiness/value). I will then propose some alternative ways for Objectivists/Subjectivists to handle these cases, but I will argue that these ways are either implausible or tacitly appeal to new, obscure normative concepts, in which case Objectivists/Subjectivists would be better off appealing to an additional sense of "ought" instead. Doing that is the only way for Objectivists (Subjectivists) to fully accommodate our feeling that, in at least some scenarios, an agent does what she ought to do even though her action is not the most appropriate thing to do in light of (her beliefs about) the facts.

3. The Subjectivist Approach

As we've seen, Subjectivists can respond to the Twin Towers case by saying that although Tom does what he ought to do when he instructs the workers to take the stairs, his action has horrible consequences—there is another action he could have performed that would have had much better consequences. Thus, Subjectivists can claim that they have the resources to say both positive and negative things about the case and thereby accommodate the intuitive tension it evokes.

However, while this way of responding to the Twin Towers case is sensible, it won't work for other, deontological versions of the case:

> Airport 1
> Two of my friends, Jesse and Andrew, need to be picked up from the airport. They are at different airports, and I can pick up only one of them. I made a promise to one of the two that I will pick him up, but I made no such promise to the other. I believe, and all my evidence suggests, that I made the promise to Jesse. However, my belief is mistaken: I actually made

[6] My talk here of the "objective ought" and the "subjective ought" is shorthand for the "objective sense of 'ought'" and the "subjective sense of 'ought,'" respectively. For ease of exposition, I will often use shorthand like this in what follows.

> the promise to Andrew. Suppose, however, that the consequences of my two options are equally good: no matter which friend I pick up, it will be a minor inconvenience to the other. In the end, I choose to pick up Jesse.

Have I done what I ought to do? This case elicits the same tension that the Twin Towers case elicits. On the one hand, I justifiably believe that I made the promise to Jesse, so perhaps I did the right thing. On the other hand, I actually made the promise to Andrew, so perhaps I did the wrong thing. However, a Subjectivist won't be able to respond to this case in the usual way. The usual way for a Subjectivist to respond would be to say that all that is intuitive is that my action is subject to both a positive and a negative overall moral evaluation, but this is consistent with Subjectivism because Subjectivists can say (positively) that when I picked up Jesse I did the right thing, but (negatively) doing so did not have the best overall consequences. However, a Subjectivist can't say this because it's simply not true that picking up Jesse fails to have the best consequences: by stipulation, the consequences of my alternatives are equally good.

There are two lines of response a Subjectivist might pursue at this point.

First, she might adopt a hardcore type of Consequentialism and deny that there is any intuitive tension in Airport 1 to resolve. The case obviously depends on promise-keeping/breaking having moral significance apart from its impact on overall consequences. A Subjectivist might simply deny this assumption.

Second, a Subjectivist might argue that I have calculated the value of the consequences incorrectly in the case and that picking up Jesse *does* in fact fail to have the best overall consequences. This is because, a Subjectivist might say, the fact that an action keeps a promise is *itself* a good consequence (apart from any satisfaction the promise-keeping brings about), and the fact that an action breaks a promise is *itself* a bad consequence (apart from any dissatisfaction the promise-breaking brings about). When we take this into account, a Subjectivist might say, we will see that picking up Andrew does in fact have better consequences than picking up Jesse since the former keeps a promise and the latter breaks a promise.

I think the second line of response is more promising than the first, and I shall therefore focus on it. Following W. D. Ross (1930), I think it's highly plausible to maintain that promise-keeping/breaking has intrinsic moral significance. Of course, this is a controversial claim, and I shall not defend it here. However, if I can establish that the tenability of Subjectivism rests on embracing a Consequentialist version of the view, that would be an interesting result in its own right.

Let's return, then, to the second line of response. While it is more plausible than the first, it doesn't stand up to scrutiny. To see this, consider a variant of the airport case:

Airport 2

The case is just like Airport 1 except that now my picking up Jesse allows him to keep a promise that *he* has made to pick up one of *his* friends. If I hadn't picked up Jesse, he would not have been able to keep this promise. Andrew, on the other hand, made no such promise. Assume that the rest of the consequences of my options are the same. That is, assume that the consequences of my two options are equally good overall.

When I pick up Jesse in this case, have I done what I should do? There's still tension to resolve in this case: on the one hand, it seems like I should pick up Jesse because I think I made a promise to him; on the other hand, it seems like I should have picked up Andrew instead because I in fact made a promise to him. The fact that picking up Andrew would have caused Jesse to be unable to keep *his* promise does not remove the inclination to say that I should, in some sense, have picked up Andrew in the case—when determining what *I* should do, the fact that my action keeps a promise *I* have made is much more important, morally speaking, than the fact that my action prevents someone else from keeping a promise *he* has made. However, a Subjectivist can't accommodate the tension generated by this case by saying that I ought to pick up Jesse but picking up Andrew has the best overall consequences since keeping a promise is itself a good consequence and breaking a promise is itself a bad consequence, for in this case, each of my options results in the keeping of one promise and the breaking of one promise: picking up Jesse causes me to break my promise and allows Jesse to keep his; picking up Andrew results in my keeping my promise and causes Jesse to break his. My alternatives are on a par with regard to the amount of promises kept, the amount of promises broken, and the value of the rest of the consequences involved.

In response, a Subjectivist might say that what Airport 2 shows is that the value of promises should be relativized to agents. Other things being equal, the fact that my action allows me to keep a promise I have made is more valuable *relative to me* than the fact that my action allows Jesse to keep his promise. But relative to Jesse, the fact that his action allows him to keep a promise is, other things being equal, more valuable than the fact that his action allows me to keep my promise. A Subjectivist might then use this concept of goodness *relative to a person* to respond to Airport 2. The Subjectivist might say that although I ought to pick up Jesse since I believe I promised him to do so, picking up Andrew is better, *relative to my perspective*, than picking up Jesse (since keeping my promise to Andrew is better, from my perspective, than allowing Jesse to keep his promise).[7]

The problem with this move is that the concept of goodness relative to a person is obscure. There is an ordinary notion of agent-relative goodness—namely, the sort of goodness that Ethical Egoists appeal to. Ethical Egoists hold

[7] Portmore (2005) appeals to this type of value.

that a person morally ought to do what is best *for her*, which I take to be equivalent to the claim that a person morally ought to do what is in her best interest to do. But the concept of goodness relative to a person that the Subjectivist is appealing to in the previous paragraph cannot be the same as the Egoist's notion of goodness for a person because the Subjectivist wants to say that my picking up Andrew is best, relative to me, but we know nothing about whether doing so is in my best interest. It might be that I'd enjoy Jesse's company more than Andrew's and so it is in my best interest to pick up Jesse after all. In that case, it still seems that I should, in some sense, have picked up Andrew since I promised him that I would do so.

So, then, when our imagined Subjectivist appeals to the concept of goodness relative to a person, what is she talking about? It's not clear.[8] Perhaps the Subjectivist can shed some light on this concept for us; however, even if she can, she will be appealing to something new, something that is not uncontroversially a part of the conceptual moral landscape. And if that's the case, then the Subjectivist may as well posit the existence of an objective ought instead. After all, as demonstrated by the passages cited at the beginning of this paper, what many people find intuitive about the cases we've been discussing is not merely that the action in question is subject to both a positive and a negative moral evaluation, but rather, that the action seems, from one perspective, to be *morally right*, but from another perspective, it seems *morally wrong*. And Subjectivists cannot accommodate this precise thought, even with the concept of goodness relative to a person. That is a significant cost. So, if Subjectivists must posit the existence of a new moral concept to handle these cases, it seems to me that the new concept they posit should be the concept of an objective ought.

Additionally, positing the existence of an objective ought would help Subjectivists make better sense of cases like the Huck Finn case:

> Huck Finn
> Huck is helping his friend, Jim, escape from slavery. As they are floating down the Mississippi River, Huck starts to feel guilty about what he's doing; he believes that it is terribly wrong to help a slave escape and that doing so could land him in Hell. After wrestling with the issue for some time, Huck decides to turn Jim in. Soon after, Huck has the perfect opportunity to do just that, but he finds that he cannot go through with it. So, he continues to help Jim escape even though he thinks he shouldn't.

When Huck decides to continue helping Jim, does he act rightly? It is extremely tempting to say that, in some sense of the term, Huck does what he ought to do. However, Huck thinks it's wrong—even a mortal sin—to help Jim escape. And

[8] Schroeder (2007) expresses similar puzzlement.

we can easily imagine that, due to Huck's upbringing and cultural position, his evidence also strongly suggests that it's wrong to do so. Because of this, it's difficult to see how a Subjectivist could maintain that Huck does what he ought to do in the case. Of course, Subjectivists might just accept that Huck acts wrongly and then try to do some damage control—they could say that while Huck acts wrongly when he helps Jim, he's praiseworthy for doing so, doing so has the best consequences, etc. But while this might minimize the damage, it won't completely eliminate it, for it will not change the fact that Subjectivists are admitting that there is no moral sense of the term according to which Huck *ought* to help Jim, and that feels like a significant cost. And my point is that if the Airport cases force Subjectivists to posit the existence of either a new agent-relative type of value or an objective moral ought, they would be wise to choose the latter, for then they can not only handle the Airport cases, but they can also say what we wanted to say about Huck all along—namely, that he ought, in some legitimate moral sense of the term, to help Jim.

I conclude that Sense-Splitting is preferable to Subjectivism. There is an intuitive tension to resolve in the cases proposed by Jackson, Oddie and Menzies, Smith, and others. The natural way for Subjectivists to respond to these cases is to try to accommodate the intuitive tension they generate by appealing to the distinction between what ought to be done and what has the best consequences. But I have argued that this strategy won't work for deontological versions of the cases that involve promise-keeping, unless a Subjectivist is willing to embrace (an implausible) form of hardcore Consequentialism or appeal to a new-fangled concept of agent-relative moral value, at which point the Subjectivist would be better off appealing to objective obligation instead. So, unless a Subjectivist has some other substitute for objective obligation that can satisfactorily deal with the cases at hand, and I don't see what that would be, then Subjectivism ought to be abandoned for Sense-Splitting.

4. The Objectivist Approach

Let's turn now to Objectivism. Recall the standard Objectivist response to the Twin Towers case. Objectivists can say that when security guard Tom tells the office workers to take the stairs, he acts wrongly since this in fact leads to their death, but he is blameless for his action since he has no reason to believe that it would have terrible results. Thus, Objectivists can insist that they can say both positive and negative things about the case and thereby accommodate the intuitive tension it generates.

The problem, though, is that this response won't work for a modified version of the case:

Twin Towers 2

The case begins just like the original case except that now the protagonist is Tim, not Tom: following the crash of an airplane into the Twin Towers,

> security guard Tim instructs the office workers to evacuate via the stairs rather than the elevator. Tim justifiably believes that the stairs are safer; however, he is mistaken: taking the stairs takes too long and the office workers die when the building collapses. Had they taken the elevator instead, they would have been fine.
>
> So far, everything is the same. But now let's add the following: unlike Tom, Tim is a terrible person. He would very much like to see the office workers die. In a normal situation, he would therefore instruct the workers to take the elevator since doing that has the best chance of leading to their demise. However, the present case is an unusual one. Tim happens to know that one of the office workers, Pam, recently had foot surgery and it will be painful for her to use the stairs. To be sure, it's still in Pam's best interest to use the stairs rather than risk using the elevator, and Tim realizes this. But he also knows that using the stairs will be quite uncomfortable for her. And *that*, let's suppose, is why Tim tells the office workers to use the stairs. Telling them to take the elevator, he believes, is too risky: they might all survive. But if they take the stairs, there's a guarantee that at least someone, namely Pam, will be harmed.

When Tim instructs the office workers to use the stairs, does he do what he ought to do? This case evokes the same tension that the original case does. On the one hand, Tim does what he thinks, and his evidence suggests, is best for the office workers on the whole. So, that pushes us to say that he did what he ought to do. On the other hand, Tim's instruction leads to the death of the workers, so that pushes us to say that he acted wrongly. Plus, Tim acts from a vile motive, so that pushes us to yet again evaluate his action negatively. So, there is tension in this case to resolve. However, Objectivists can't resolve this tension in the usual way. The usual way for Objectivists to resolve the tension would be to say (negatively) that Tim does the wrong thing, but (positively) he is blameless for doing so. But Objectivists *can't* say this because Tim *isn't* blameless for his action. Since he acts from a horrendous motive, he is presumably blameworthy for his behavior.

How else might Objectivists handle this case? It seems to me that Objectivists' best recourse is to appeal in some way to virtue. Of course, it would be implausible to suggest that Tim is a virtuous person overall—indeed, it's a stipulation of the case that he is a vicious person. However, even bad people can do virtuous things from time to time (and even good people can do vicious things on occasion). And perhaps this is what is happening in Twin Towers 2: while Tim is a bad person, acts wrongly, and is blameworthy for his action, he nonetheless acts virtuously. If this is correct, then Objectivists can argue that they can accommodate the intuitive tension of the case without having to appeal to a second overall moral ought.

However, for this to work, Tim has to be exhibiting a moral virtue when he tells the workers to take the stairs. But what virtue could that be? As far as I can see, the only virtue that might be relevant is the virtue of conscientiousness. But does Tim act conscientiously in the case? Perhaps. Although there is debate about what conscientious behavior involves—more on that in a moment—at first glance, it seems plausible to hold that an agent acts conscientiously when she does what she thinks she morally ought to do, or, at least, what she thinks she is morally permitted to do. But in Twin Towers 2, we're told that Tim recognizes that it's best, on the whole, for the workers to take the stairs—even for Pam. So, it's plausible to assume that Tim believes he's doing the right thing.

While appealing to the virtue of conscientiousness may help Objectivists respond to Twin Towers 2, there are other cases where such an appeal may be more difficult. Consider, for instance, Frank Jackson's (1991) Dr. Jill case:

> Dr. Jill
> Dr. Jill has a patient, John, who is suffering from a minor skin ailment. Jill can give him one of three drugs: A, B, or C. She knows that drug B is an imperfect, but still pretty good, treatment. She also knows that one of A or C will perfectly cure John and the other will kill him, but she doesn't know which is which—her evidence is equally divided on the matter (though as a matter of fact, drug A is the perfect cure and drug C is the killer). Jill decides to give John drug B, the safe drug.[9]

Did Jill do what she ought to do? On the one hand, given what Jill knew about the available drugs, it may seem like she did in fact do what she should have done. But, on the other hand, it may seem like Jill should have given drug A, not B, since A would have cured John perfectly. Of course, Sense-Splitters can easily accommodate these conflicting thoughts: they can maintain that when Jill gave drug B, she did what she subjectively ought to do, but she did not do what she objectively ought to do. Moreover, Objectivists can also handle this case, and in the usual way. They can say that what is intuitive is simply that Jill's action is subject to both a positive and a negative moral evaluation, which is consistent with Objectivism because Objectivists can say—negatively—that Jill acts wrongly, but—positively—she is blameless for her behavior.

---

[9] Regan (1980: 264-65), J. Ross (2006: 174), and Parfit (2011: 159-60) develop similar cases. Note that the Dr. Jill case is importantly different from the cases I have discussed so far. All of the cases discussed up to this point feature protagonists who have misleading or false information. However, Dr. Jill doesn't have misleading information or any false beliefs. Instead, she has incomplete information—she doesn't know which drug is the killer and which is the cure. In section 5.2, we will see how incomplete-information cases might be thought to pose their own special problem for Sense-Splitters by forcing them to adopt an annoying profusion of oughts.

However, we can easily add details to the case so that Jill is not blameless for giving B. Suppose, for instance, that Jill hates John (for no good reason—she doesn't like how he looks, perhaps). She would love to kill him: if she knew which drug would kill him, she would give him that drug. However, Jill would hate even more to cure John's skin ailment completely. So, she decides to give him drug B because she's risk averse and wants to ensure that he continues to suffer at least partially. Now, it no longer seems that Jill is blameless for her action, and yet it still seems that she has, in some sense, done the right thing (though for the wrong reason!).

At this point, Objectivists might again try to appeal to virtue. They might say that when Jill gives B, she acts wrongly (since it's not the perfect cure), and she's blameworthy for doing so (since her motives are vile), but she exhibits the virtue of conscientiousness. But is it plausible to hold that Jill acts conscientiously in this case? The answer crucially depends on what Jill believes about her options. Suppose, for instance, that Jill is an objective Consequentialist: she believes that she ought to do the best that she can. Because of this, let's suppose that she believes that it's *wrong* to give drug B since she knows that giving this drug will *not* have the best overall consequences. If Jill is like this, then it's difficult to see how she could be acting conscientiously when she gives drug B. Plausibly, conscientious action precludes doing what one thinks is morally wrong.

While some philosophers agree with this plausible contention about conscientious action,[10] not all do. For instance, Peter Graham (2010) suggests that conscientious agents have two overarching goals—doing what they ought to do, and avoiding wrongdoing—and conscientious agents do what they think best satisfies these goals taken together. However, even as an objective Consequentialist, Jill might think that giving B best satisfies these goals taken together. Jill may judge that giving B does not at all satisfy the first goal of doing what she ought to do. She may also believe that it doesn't satisfy the second goal of avoiding wrongdoing either, but it does *pretty well* with respect to that goal since it is at least not seriously wrong. On the other hand, Jill may believe that one of A or C satisfies both goals perfectly (since it is obligatory and not wrong), but the other fails to satisfy each goal spectacularly (since it is not obligatory and is seriously wrong). All told, then, Jill might judge that giving B is the best way of balancing the twin goals of the conscientious person, and if that's the case, Graham suggests, she acts conscientiously when she gives B.

However, I wonder what it means to "best balance" the twin goals of doing what one ought to do and avoiding wrongdoing. Graham may be appealing to a new normative concept here, which would be bad if one of the aims of Objectivism is to handle cases like Twin Towers and Dr. Jill without positing the existence of

[10] See, esp., Zimmerman (2008: chap. 1; 2014: chap. 4).

new normative concepts. If the notion of best balancing the twin goals of the conscientious person is in fact a new addition to the conceptual moral landscape, then for reasons similar to the ones given above, I'd say that Objectivists are better off positing the existence of subjective rightness instead. Only by doing that can Objectivists accommodate the intuition many people have that the agents in the cases we've been discussing act both rightly and wrongly, morally speaking. Objectivists cannot accommodate this precise thought. At most, they can accommodate the sense that there are both positive and negative moral things to say about the agents involved. But I suspect that many will feel that this simply doesn't go far enough. This is especially so for Dr. Jill. It is extremely tempting to say that when Jill gives the safe drug, she does what she morally ought to do. Objectivists, it appears, must deny this. They can perhaps minimize the damage by saying that although Jill doesn't act rightly, in any sense of the term, she acts conscientiously. But this isn't fully satisfying. It sure seems like Jill is doing the right thing, in some legitimate sense of the term. Objectivists cannot accommodate this exact feeling.

I conclude that Sense-Splitting is superior to Objectivism. Objectivists can handle many of the cases we've encountered by appealing to the distinction between rightness/wrongness and praiseworthiness/blameworthiness. However, I have argued that appealing to this distinction will not handle all test cases. Objectivists might try to appeal to virtue/vice to handle the remaining cases. But I have argued that the success of this strategy depends on accepting some controversial claims about virtue. And even if Objectivists are willing to accept those claims, there is still a cost for the view since no matter what acrobatics Objectivists do to accommodate the test cases, they can never fully satisfy the powerful urge to say that the agents in these cases act rightly. The only way to fully accommodate this urge is to say that the agents do in fact act rightly, in some sense of the term. And the only way for Objectivists to accommodate *that* is to embrace Sense-Splitting. So, unless Objectivists have some other substitute for subjective obligation that can handle the cases at hand, and I don't see what that would be, then Objectivism should be abandoned for Sense-Splitting.

5. Objections

Having now explained why we should accept Sense-Splitting instead of Objectivism or Subjectivism, I will defend Sense-Splitting from two important objections.

5.1. Which "Ought" Is More Important?

Sense-Splitters hold that there is an objective and a subjective sense of "ought," but this leads to natural questions when we think about cases where a person ought—in the objective sense—to do one thing, but she ought—in the subjective sense—to do a different thing. In such cases, which ought is "more important"? Which ought "ought" the person follow? Which ought "really

counts"? Several philosophers have raised these types of questions about Sense-Splitting.[11]

In response, note first that these remarks are more of a request for further information than a straightforward objection. What they make clear is that Sense-Splitting leads to natural questions that Sense-Splitters should be prepared to answer. Of course, Sense-Splitters may have no plausible way of answering those questions, and *that* would be a problem. But it remains to be seen whether this is the case.

However, as I will now show, there are several promising ways for Sense-Splitters to answer these questions. I won't try to defend one of these ways in particular. My goal here is more modest. I merely want to demonstrate that Sense-Splitters have some plausible avenues for answering the questions. These questions are therefore not suggestive of a deep problem for the view.

First, let's narrow the *i*mportant *q*uestion that I'll be trying to answer:

$IQ_{\text{sense-splitter}}$: What ought I to do when my subjective obligation conflicts with my objective obligation?

When one asks whether it's the subjective ought or the objective ought that is "more important" or "really counts," I suspect that one is really just asking $IQ_{\text{sense-splitter}}$. I will therefore focus on it.

Notice, however, that *everyone*, even Objectivists and Subjectivists, must grapple with a version of $IQ_{\text{sense-splitter}}$. After all, virtually everyone admits that there are many types of overall oughts: the prudential ought, the legal ought, the aesthetic ought, the moral ought, etc. In light of this, it is natural to wonder, which of these oughts is the most important? Which one trumps the others? In situations where they diverge, which one ought we to follow? Everyone, therefore, must confront an important question like $IQ_{\text{sense-splitter}}$:

IQ_{everyone}: What ought I to do when there is conflict between my moral, prudential, legal, aesthetic, etc. obligations?

I will now discuss some plausible ways of responding to IQ_{everyone} and then apply them to $IQ_{\text{sense-splitter}}$.

One way of responding to IQ_{everyone} is to say that it is incoherent. Indeed, the question is *prima facie* puzzling. When we ask what ought I to do when there is a conflict between my moral, prudential, legal, etc., obligations, which ought do we have in mind? We surely don't have the moral ought in mind; otherwise, the

[11] See, e.g., W. D. Ross (1939: 147) and Zimmerman (2008: 7).

question would be silly: obviously, I *morally* ought to follow my moral obligation when there is a conflict between my moral, prudential, legal, etc. obligations. Similarly, when we ask what ought I to do when there is a conflict between my moral, prudential, legal, etc., obligations, we surely don't have the prudential ought in mind, because once again, the question would be silly. It may therefore seem that $IQ_{everyone}$ is meaningless. The ought it is asking about can't be the moral ought or the prudential ought. For similar reasons, it can't be the legal ought, the aesthetic ought, or, it may seem, *any ought whatsoever*. $IQ_{everyone}$ therefore looks hopelessly confused.

However, one might argue that there *is* an ought that would make $IQ_{everyone}$ sensible. One might agree that the ought at issue in $IQ_{everyone}$ is not the moral ought, the prudential ought, the legal ought, or any ought that is relativized to a specific normative framework. Instead, one might say, the ought at issue is one that "transcends" all of the individual normative perspectives. The ought at issue is the "all-things-considered" ought, or the "just-plain" ought.

What is the just-plain ought? As I suggested, it's not relativized to any particular normative perspective, and it somehow transcends the oughts that are relativized to a perspective. We might also add that the just-plain ought is somehow "superior to" or "weightier than" these other oughts. But this is all pretty vague. What else can we say about the concept? Can it be made clearer?

These are difficult questions. Some philosophers are deeply puzzled by the just-plain ought and think that the concept is incoherent.[12] Others, though, seem to be able to grasp the concept.[13] For my part, I find that I teeter between each side of the dispute. On the one hand, it seems to me that when there is a conflict between our moral, prudential, legal, etc. obligations, there is some real sense in which I ought to follow my moral obligation, where the "ought" at issue is not the moral ought, the prudential ought, or any other framework-relative ought. This makes me think that I have some grasp of the concept of a just-plain ought. But that grasp is contingent on there being a tight connection between the moral ought and the just-plain ought: it seems to me that someone just-plain ought to do something iff the person morally ought (in some sense) to do it. I have few intuitions about the just-plain ought beyond this. But not everyone shares this view. Some hold that the just-plain ought and the moral ought can come apart.[14] But if the just-plain ought is supposed to be an ought that can come apart from

[12] See, e.g., R. Feldman (2000), Copp (1997), and Baker (forthcoming).

[13] See, e.g., McLeod (2001), F. Feldman (1986), Darwall (1997), and Chang (2004).

[14] See, e.g., Portmore (2008).

the moral ought, then perhaps I don't understand the concept after all. So, I feel the pull of each side of the dispute.

In any case, there are two broad strategies for responding to $IQ_{everyone}$. First, one can claim that the question is not worrisome because it is incoherent—it depends on the confused notion of a just-plain ought. Second, one can claim that the question is coherent since the concept of a just-plain ought makes sense, and one can then answer $IQ_{everyone}$ by appealing to a theory of just-plain obligation. It's beyond the scope of this paper to determine whether the just-plain ought is coherent and, if it is, what the correct theory of it looks like. All I'm trying to establish is that there are some promising ways of answering $IQ_{everyone}$; the fact that the question emerges when we admit that there is a plurality of perspective-relative oughts gives us no good reason to doubt that a plurality of these oughts exists.

What I have said about $IQ_{everyone}$ applies straightforwardly to $IQ_{sense\text{-}splitter}$. When one asks, what ought I to do when my objective obligation conflicts with my subjective obligation, the ought at issue is surely not the objective ought or the subjective ought—otherwise the question would be silly. Rather, the ought at issue is some third ought. And what is this ought? A natural answer, given the foregoing, is that the ought at issue is the just-plain ought. $IQ_{sense\text{-}splitter}$ would then be asking, what just-plain ought I to do when my objective obligation conflicts with my subjective obligation? If that is what $IQ_{sense\text{-}splitter}$ is asking, then there are, we've seen, two broad strategies for answering it. Either one can claim that the question is incoherent because the concept of a just-plain ought is incoherent, or one can claim that the question is coherent and offer a theory about what I just-plain ought to do in these circumstances. Each strategy is, I think, a reasonable way of addressing the question.

However, one might argue that what I've said about $IQ_{everyone}$ doesn't apply straightforwardly to $IQ_{sense\text{-}splitter}$, for there are reasons for thinking that the just-plain ought, even if it is coherent, cannot incorporate both objective and subjective factors and so cannot settle the question about what I should do when my objective obligation conflicts with my subjective obligation. One might think that the just-plain ought cannot incorporate both objective and subjective factors if we adopt a semantics of "ought" along the lines of Kratzer (2002).[15] On Kratzer's view, "ought"-statements involve two contextually-determined parameters: (1) a "modal base" that selects the possible worlds that are available for assessment, and (2) an "ordering source" that ranks the worlds selected by the modal base. But, one might argue, while it's easy to see how an ordering source could rank worlds on the basis of both objective (factual) and subjective (epistemic) criteria, it's not clear that a modal base could select worlds on the basis of both of these factors. To illustrate, it may be that the modal base of the objective ought selects all and only those

[15] Thanks to an anonymous referee for suggesting this line of thought to me.

worlds that the agent can *in fact* bring about, whereas the modal base of the subjective ought selects all and only those worlds that the agent *believes* she can bring about. Since these two sets of worlds can be different, it is difficult to see what the modal base of an ought that incorporates both the objective and the subjective ought would look like.

Note, however, that some philosophers believe that the subjective "ought" implies "can" (not merely "*believes* can") in the same way that the objective "ought" does.[16] For such philosophers, the modal base selected by the subjective ought will presumably be the same as the one selected by the objective ought—both will select all and only those worlds the agent can in fact bring about. But even if the modal bases are different, this will threaten the coherence of a just-plain ought that governs both the objective ought and the subjective ought only if we assume that a just-plain ought of this sort must use both objective and subjective criteria to select its modal base. But why must we think that? As I've suggested above, it might be that the just-plain ought is very closely connected to the subjective ought—the two are coextensive; factors that determine whether someone just-plain ought to do something just are the factors that determine whether someone subjectively ought to do it. If the just-plain ought were like this, then the modal base of the just-plain ought would presumably be selected via the same criteria as the subjective ought. I don't see why a view like this should be ruled out. Of course, it would be ruled out if we held that the criteria used by the just-plain ought to select its modal base must be identical to that of every ought it governs. But this assumption needs defending, and I don't see what a defense of it would look like.

Let me also point out that the line of reasoning we are considering would, if sound, likely provide grounds for the wholesale rejection of the just-plain ought. The just-plain ought is supposed to govern all of the perspective-relative oughts—the moral ought, the prudential ought, the legal ought, the ought of rationality, etc. But even if there is only one moral ought, it's unlikely that its modal base is selected in the same way with the same objective or subjective criteria as that of every other perspective-relative non-moral ought. Moreover, as some have pointed out, many of the non-moral oughts (such as the prudential ought and the ought of rationality) appear susceptible to the objective/subjective distinction, just as the moral ought is.[17] So, if we should doubt that there could be a just-plain ought that governs both the objective moral ought and the subjective moral ought because the just-plain ought would need to have a modal base selected on the basis of both objective and subjective factors, then I suspect that we'll also have to doubt, for the same reason, that there could be a just-plain ought that governs the perspective-relative oughts

[16] See, esp., Zimmerman (2008: sec. 3.5; 2014: chap. 1). For a contrasting view, see Hedden (2012).

[17] See, e.g., Wedgewood (2007: 118; 2016: 147).

more generally. But doubting this is tantamount to doubting that there is any such thing as the just-plain ought at all, since a defining feature of the just-plain ought is that it governs the perspective-relative oughts.

Thus far, I have assumed that the ought at issue in $IQ_{\text{sense-splitter}}$ is the just-plain ought. If that is correct, then I've argued that the question can be answered in the same general way as IQ_{everyone}. However, even if this is incorrect, there is another way of answering $IQ_{\text{sense-splitter}}$ that won't also apply to IQ_{everyone}: we could interpret the ought in $IQ_{\text{sense-splitter}}$ not as the just-plain ought, but as an ought that is relativized to a normative framework. So, when we ask what ought I to do when my objective obligation conflicts with my subjective obligation, we might be asking what I *prudentially* ought to do in these circumstances, or what I *aesthetically* ought to do, or what I *legally* ought to do. More likely, we might be asking about a *means/end* ought. That is, we might have in mind the sort of ought according to which it's the case that I ought to bring my umbrella on the assumption that I want to stay dry, and I ought to arrive at the train station ten minutes early on the assumption that I want to avoid missing the train. On this ought, what one ought to do depends on what goals or ends one is trying to achieve: if one has a goal, and doing something is sufficient for achieving that goal, then one ought to do that thing. If this type of ought is at issue in $IQ_{\text{sense-splitter}}$, then the meaning of the question will depend on what goals the questioner has in mind. Thus, the question might amount to: what ought I to do when my objective and subjective obligations conflict *on the assumption that I want to be blameless*? Or it might mean: what ought I to do when my objective and subjective obligations conflict *on the assumption that I want to avoid feeling guilty*? Or perhaps: what ought I to do when my objective and subjective obligations conflict *on the assumption that I want to do what a morally conscientious person would do*? Of course, if $IQ_{\text{sense-splitter}}$ amounts to one of these things, we still need to decide how to answer it. I suspect that the answer to all of these questions is the same: I ought, in the means/end sense at issue, to follow my subjective obligation. But I won't argue for that here. I'm merely suggesting another potentially promising way of handling $IQ_{\text{sense-splitter}}$, namely by interpreting the ought at issue as a perspective-relative ought.

This concludes my discussion of $IQ_{\text{sense-splitter}}$. What I've said has been more of a roadmap than an argument. My aim has not been to defend any particular answer to the question, but, rather, to establish that there are several reasonable ways of addressing it. More generally, my goal has been to demonstrate that the fact that $IQ_{\text{sense-splitter}}$ naturally arises for anyone who adopts Sense-Splitting gives us little reason to doubt the view. It is not a question with no plausible answer. Rather, it can be addressed in several promising ways.

5.2. An Annoying Profusion of Oughts

Some have suggested that Sense-Splitting is problematic because it leads to an "annoying profusion of oughts," as Jackson (1991: 471) puts it.[18] The idea is that if one is going to draw a distinction between what a person ought to do in light of the facts of her situation (the objective ought) and what the person ought to do in light of her beliefs about the facts of her situation (the subjective ought), then one will be committed to accepting many other oughts as well: an ought that corresponds to what the person ought to do in light of *my* beliefs about the facts of the situation, an ought that corresponds to what the person ought to do in light of *your* beliefs about the facts of the situation, an ought that corresponds to what the person ought to do in light of *God's* beliefs about the facts of the situation, and so forth. But, the objection goes, it is problematic to accept so many moral oughts.

This line of objection is unpersuasive. First, even if Sense-Splitters are committed to a profusion of oughts, this isn't a problem if Sense-Splitters can rank them. If we can say, for instance, that the subjective moral ought is "more important" than the other moral oughts, in the sense that we just-plain ought to follow the subjective ought when it conflicts with the other oughts, then this makes the profusion of oughts much more palatable. Of course, there will be debate about how to rank the profusion of oughts in terms of importance. But the point is that if Sense-Splitters can impose a hierarchy of importance on the profusion, this makes it much less annoying. And if Sense-Splitters are willing to countenance the notion of a just-plain ought, then it seems that they have adequate resources to impose such a hierarchy.

Second, and more importantly, it's not clear that Sense-Splitters need to accept the annoying profusion of oughts in the first place. Sense-Splitters hold that there is an ought that's relativized to the facts of the agent's situation. They also hold that there is an ought that's relativized to the agent's mental state. But these claims, on their own, do not force Sense-Splitters to accept that there are additional oughts that are relativized to the mental states of others. More than just the bare commitments of Sense-Splitting are needed to compel its proponents to accept Jackson's annoying profusion of oughts.

However, Sense-Splitters will have to accept an annoying profusion of oughts if the same line of reasoning they use to motivate the distinction between the objective ought and the subjective ought can be used equally well to motivate the existence of a profusion of additional oughts. Recall what that motivation was: Sense-Splitters motivate the distinction between the objective ought and the subjective ought by appealing to cases where part of us wants to say that an agent ought to do one thing (when we look at the facts), but another part of us wants to say that the agent ought to do a different thing (when we look at the agent's mental

[18] Cf. Kolodny and MacFarlane (2010: 121).

state). So, for instance, it is tempting to say that Tom ought to tell the office workers to take the stairs (since he thinks that doing so is safest), but it is also tempting to say that Tom ought to tell the workers to take the elevator (since doing so will in fact save their lives). This provides motivation for drawing a distinction between an objective ought and a subjective ought. So, Sense-Splitters will be committed to accepting the other oughts that Jackson mentions if a clear case can be found where part of us wants to say that an agent ought (in the objective sense) to do one thing, another part of us wants to say that the agent ought (in the subjective sense) to do another thing, and still another part of us wants to say that the agent ought (in the other senses Jackson mentions) to do other things. But what would such a case look like?

One of the most compelling cases of this type is developed by Kolodny and MacFarlane (2010: 121):

> Dr. Jill's Advisor[19]
> Jill goes to Jack for advice about what to do. She knows that Jack has better evidence than she does but that he does not know which of the drugs is the perfect cure. She asks Jack, "what ought I to do in my circumstances?" Let's suppose that Jack is sure that the partial cure is C (not B, as Jill believes), but he's unsure which of A or B is the perfect cure and which is the killer. In light of this, Jack tells Jill, "you ought to give drug C."

Intuitively, when Jill asks Jack what she ought to do, her question makes sense. Moreover, when Jack tells Jill that she ought to give C, he seems to speak truly. But what sense of "ought" is at issue in this conversation? Presumably, Jill doesn't have the objective ought in mind in her question since she knows that Jack doesn't know which of the drugs is the perfect cure and thus doesn't know what she objectively ought to do. Jill also presumably doesn't have the subjective ought in mind either; we can imagine that she already knows what she subjectively ought to do: give drug B. It seems, then, that Jill has in mind a third ought, one that's different from the objective ought and the subjective ought. The ought that she has in mind seems to be one that is relativized to *Jack's* evidence—after all, Jill goes to Jack for advice precisely because he has better evidence than she does. It seems, then, that the above conversation motivates the existence of an ought that is relativized to Jack's evidence. But it's easy to see how we can also use similar cases to motivate the existence of an ought that's relativized to *my* evidence, an ought that's relativized to *your* evidence, and so on. So, one might think that Sense-Splitters are committed to an annoying profusion of oughts after all.

In response, note first that if this case is a problem for Sense-Splitters, it's also a problem for Objectivists and Subjectivists. As we've seen, it appears that Jill

[19] I have modified this case slightly for my purposes.

is not asking Jack either about what she objectively ought to do or about what she subjectively ought to do. So, even Objectivists and Subjectivists have to grapple with what Jill means when she asks Jack what she ought to do. The case does not pose a special problem for Sense-Splitting.

Setting that aside, how should a Sense-Splitter accommodate the alleged fact that Jill's question makes sense and Jack's answer is true? Can a Sense-Splitter accommodate these facts without positing the existence of a third ought that is relativized to Jack's evidence?

I believe so. There are several ways forward. First, a Sense-Splitter can say that the ought at issue in the case is indeed an ought that is distinct from the objective ought and the subjective ought, but it's not an ought that is relativized to Jack's evidence; rather, it's the just-plain ought. This move obviously assumes that the concept of the just-plain ought is coherent and that it's plausible to hold that Jill just-plain ought to give drug C. Each of these assumptions (particularly the latter) can be questioned. Still, depending on Sense-Splitters' views about the nature of the just-plain ought and its implications in particular cases, it remains a notable option for them to try to make sense of Jack and Jill's dialogue by appealing to the just-plain ought.

Another (perhaps more promising) option would be for a Sense-Splitter to try to make sense of the dialogue by appealing to a more familiar ought, such as a means/end ought. For instance, perhaps when Jill asks Jack what she ought to do, she's asking him what she ought to do (in the means/end sense) on the assumption that she wants to act as Jack would act in her circumstances. This question makes sense. And, assuming that Jack would give C in Jill's circumstances and that he has the same means/end ought in mind in his response, then his reply (that she ought to give C) is true.

Alternatively, a Sense-Splitter might say that the ought at issue in Jack and Jill's conversation is simply the subjective ought. But how could that be, given that I am assuming that Jill already knows what she subjectively ought to do? Well, maybe when she asks Jack what she ought to do, she's asking him about a *conditionalized* subjective ought. Maybe, for instance, she's asking him what she ought to do (subjectively) *if she were in his shoes*. Or perhaps she's asking him what she ought to do (subjectively) *if she had the same evidence as Jack*. Another option: perhaps when Jill asks Jack what she should do, she isn't asking him about a conditionalized subjective ought, but rather, she's asking him what she ought (subjectively, unconditionally) to do *at some later time*. In particular, perhaps she is asking Jack what she ought to do *later, after Jack informs her of his superior evidence*.[20] All of these questions make sense. And, assuming that Jack has one of these

[20] Cf. Kiesewetter (2011) and Zimmerman (2014: 82-87).

varieties of the subjective ought in mind in his response, then his reply—that she should give C—looks true.

So, there are several ways for a Sense-Splitter to accommodate the intuitive data of Dr. Jill's Advisor without appealing to an ought that is relativized to Jack's evidence. The case therefore does not seem to me to provide a strong reason for saddling Sense-Splitters with Jackson's annoying profusion. However, one might argue that there are other, more complex versions of the case that are more difficult for Sense-Splitters to accommodate. Consider, for instance, the following case, developed by Mark Schroeder (forthcoming):

> Four Envelopes[21]
> Xiao gets to choose one of four envelopes. One of the envelopes contains $4000, one contains $5000, one contains $6000, and one contains nothing. Xiao knows this and that envelope A contains $4000, but she doesn't know what amounts the other envelopes contain.
>
> Ying knows all of the information that Xiao knows (including what Xiao knows), but Ying also knows that envelope B contains $5000. He doesn't know, however, whether the $6000 is in envelope C or D.
>
> Zach knows all of the information that Ying knows (including what Ying knows), but Zack also knows that envelope C contains $6000.
>
> Now consider the following monologues:
> Xiao thinks to herself: "I ought to take envelope A."
> Ying thinks to himself: "Xiao ought to take envelope B."
> Zach thinks to himself: "Xiao ought to take envelope C."

Schroeder notes that each of these monologues makes sense on its own. Furthermore, it would make sense for the speaker of each of these sentences to say it to Xiao if she were to ask for advice.

Of course, Sense-Splitters have no problem handling Xiao's and Zach's statements: Sense-Splitters can say that Xiao's utterance concerns the subjective ought and Zach's utterance concerns the objective ought. But it may seem that Ying's utterance poses a problem for Sense-Splitters, for Ying presumably doesn't have in mind the subjective ought since he knows that Xiao subjectively ought to

[21] Thanks to an anonymous referee for drawing my attention to this case. Note that while this case can be used to pose a problem for Sense-Splitting, that is not why Schroeder discusses it. He discusses it because he wants to defend the idea that oughts are determined by the balance of reasons, and he thinks that Four Envelopes poses a potential problem for such a view. Schroeder ultimately argues that the problem can be solved by embracing a certain type of expressivism about epistemic expressions.

take envelope A, nor does he seem to have in mind the objective ought since he knows that Xiao objectively ought to take either envelope C or envelope D. So, it looks like Ying is appealing to a third ought that is relativized neither to the facts nor to Xiao's mental state, but rather, to Ying's mental state. Furthermore, it is easy to see how we can expand this case by adding more envelopes and more speakers in order to motivate the existence of a plethora of oughts, each corresponding to a different person's mental state. It may appear, then, that Sense-Splitters are again forced to accept the implausible claim that there is an annoying profusion of oughts.

In fact, I don't think the problem this case raises is substantially different from the problem raised by Dr. Jill's Advisor, and I would respond to it in much the same way. First, it's worth stressing that like Dr. Jill's Advisor, Four Envelopes poses a problem for Objectivists and Subjectivists as well since Ying's utterance doesn't appear to concern either the objective ought or the subjective ought. Objectivists and Subjectivists therefore also owe us an account of Ying's statement.

More importantly, Sense-Splitters have the resources to make sense of Ying's claim without accepting a multitude of additional oughts. One way of doing this is by accepting the proposal advocated by Mark Schroeder. He argues that when Ying claims that Xiao ought to take envelope B, he may in fact have the objective ought in mind. This is because, according to Schroeder, objective oughts are determined by the balance of objective reasons, but objective reasons are, in an important way, information sensitive: what a person has objective reason to do depends on the speaker's evidence. So, when Ying says that Xiao ought to take envelope B, he may very well have the objective ought in mind, and his claim might very well be true since the objective reasons for Xiao to act, as communicated by Ying, will include the fact that the expected value of Xiao's taking envelope B is $5000, which is greater than the expected value of each of Xiao's other options.

Discussion of Schroeder's intriguing proposal is beyond the scope of this essay. Let me note, however, that even if it is implausible, this won't necessarily doom Sense-Splitting, for there are other ways for Sense-Splitters to accommodate the intuitive data in Four Envelopes. As before, Sense-Splitters can maintain that when Ying says that Xiao ought to take envelope B, he means that she ought (in the means/end sense) to take B on the assumption that she wants to act as Ying would in the situation, or that she ought (subjectively) to take B if she had all of the information that Ying has. Furthermore, if, in response to a request for advice, Ying tells Xiao that she ought to take B, Ying may be telling Xiao not what she ought (subjectively) to do *now*, prior to their conversation, but what she ought (subjectively) to do *later*, after their conversation is over and Xiao's evidential state has thereby been improved. All of these options are plausible interpretations of Ying's statement, and all of them preserve its truth.

It might be objected, though, that Ying's claim that Xiao ought to take envelope B seems true even if Ying has in mind a *moral* ought (as opposed to a means/end ought), an *unconditional* ought (as opposed to what Xiao ought to do if some set of conditions obtains), and an *immediate* ought (as opposed to what Xiao ought to do at some future time). But is it really so intuitive that Xiao morally ought to take envelope B in her current evidential situation, prior to discussing her options with a better-informed advisor? Recall that her evidence suggests that envelope B is just as likely to contain nothing as it is to contain $5000 or $6000, whereas A is guaranteed to contain $4000. Given that, it strikes me that she simply shouldn't take B, at least now, in her current position—doing so is just too risky and will not result in her getting the maximum prize (in envelope D). Of course, if Ying were to *talk* to Xiao, then it might be that he should *tell* her to take B, and after that, it might be that she actually should take B. But that's not the situation we are imagining. The situation we are imagining is one where Xiao has not discussed her situation with Ying or anyone else. At *that* time, in *that* particular situation, it seems to me, it would be a mistake for Ying to think that Xiao morally ought to take B. If Ying is concerned, in his monologue, to soliloquize the truth, he should say either that Xiao ought to take A, or that Xiao ought to take D. The fact that Ying's evidence tells him that the second-best prize is elsewhere seems to me to be irrelevant to the moral status of Xiao's options in the circumstances in which she presently finds herself.[22]

In sum, the annoying-profusion-of-oughts problem is not particularly worrisome for Sense-Splitters. Sense-Splitters needn't accept an annoying profusion of oughts. And even if they do, this won't be problematic if they can arrange the oughts in a hierarchy of importance, which they have the resources to accomplish if they countenance a just-plain ought. The two primary objections to Sense-Splitting are therefore unconvincing, and, we've seen, there is good reason to favor Sense-Splitting over its competitors. We should all be Sense-Splitters.[23]

Department of Philosophy
University of New Hampshire

---

[22] I have been assuming that Xiao has moral duties to herself, an assumption that can be challenged (see, e.g., Ross [1939]). If Xiao has no moral duties to herself, then she is arguably morally permitted to choose any of the four envelopes. We can, however, remove this complication by making Xiao's choices other-directed; for instance, we can stipulate that the money in the chosen envelope will go to charity. My conclusions about the case remain unchanged with this added stipulation.

[23] Many people gave me valuable feedback on drafts of this paper. Thanks especially to Fred Feldman, Luis Oliveira, Scott Hill, and an anonymous referee for this journal.

Works Cited

Baker, Derek. Forthcoming. "Skepticism about Ought *Simpliciter*." In *Oxford Studies of Metaethics*, 13th ed., edited by Russ Shafer-Landau. Oxford: Oxford University Press.

Brandt, Richard B. 1959. *Ethical Theory: The Problems of Normative and Critical Ethics*. Englewood Cliffs, NJ: Prentice-Hall.

Chang, Ruth. 2004. "All Things Considered." *Philosophical Perspectives* 18: 1-22.

Copp, David. 1997. "The Ring of Gyges: Overridingness and the Unity of Reason." *Social Philosophy and Policy* 14: 86-101.

Darwall, Stephen. 1997. "Learning from Frankena: A Philosophical Remembrance." *Ethics* 107: 685-705.

Dorsey, Dale. 2012. "Objective Morality, Subjective Morality and the Explanatory Question." *Journal of Ethics and Social Philosophy* 6 (3): 1-24.

Ewing, A. C. 1948. *The Definition of Good*. London: Routledge and Kegan Paul.

Feldman, Fred. 1986. *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht, Holland: D. Reidel.

———. 2012. "True and Useful: On the Structure of a Two Level Normative Theory." *Utilitas* 24: 151-71.

Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60: 667-95.

Gibbard, Allan. 2005. "Truth and Correct Belief." *Philosophical Issues* 15: 338-50.

Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121: 88-115.

Hedden, Brian. 2012. "Options and the Subjective Ought." *Philosophical Studies* 158: 343-60.

Howard-Snyder, Frances. 2005. "It's the Thought that Counts." *Utilitas* 17: 265-81.

Jackson, Frank. 1986. "A Probabilistic Approach to Moral Responsibility." In *Logic, Methodology, and Philosophy of Science VII*, edited by Ruth Barcan Marcus, Georg J. W. Dorn, and Paul Weingartner, 351-65. North Holland: Elsevier Science Publishers.

———. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101: 461-82.

Kiesewetter, Benjamin. 2011. "'Ought' and the Perspective of the Agent." *Journal of Ethics and Social Philosophy* 5(3): 1-24.

Kolodny, Niko, and John MacFarlane. (2010). "Ifs and Oughts." *Journal of Philosophy* 107: 115-43.

Kratzer, Angelika. 2002. "The Notional Category of Modality." In *Formal Semantics: The Essential Readings*, edited by Paul Portner and Barbara H. Partee, 289-323. Oxford: Blackwell.

Mason, Elinor. 2013. "Objectivism and Prospectivism about Rightness." *Journal of Ethics and Social Philosophy* 7(2): 1-21.

McLeod, Owen. 2001. "Just Plain 'Ought.'" *The Journal of Ethics* 5: 269-91.

Moore, G. E. (1903) 1993. *Principia Ethica*. Rev. ed. Edited by Thomas Baldwin. Cambridge: Cambridge University Press.
———. (1912) 2005. *Ethics*. Edited by W. H. Shaw. Oxford: Clarendon Press.
Oddie, Graham, and Peter Menzies. 1992. "An Objectivist's Guide to Subjective Value." *Ethics* 102: 512-33.
Parfit, Derek. 2011. *On What Matters*. Vol. 1. Edited and introduced by Samuel Scheffler. Oxford: Oxford University Press.
Pollock, John. 1979. "A Plethora of Epistemological Theories." In *Knowledge and Justification*, edited by G. S. Pappas, 93-113. Dordrecht, Holland: D. Reidel.
Portmore, Douglas W. 2005. "Combining Teleological Ethics with Evaluator Relativism: A Promising Result. *Pacific Philosophical Quarterly* 86: 95-113.
———. 2008. "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11: 369-88.
———. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.
Prichard, H. A. (1932) 2002. "Duty and Ignorance of Fact." In *Moral Writings*, edited by Jim MacAdam, 85-110. Oxford: Oxford University Press.
Regan, Donald. 1980. *Utilitarianism and Co-Operation*. Oxford: Clarendon Press.
Ross, Jacob. 2006. *Acceptance and Practical Reason*. Ph.D. Thesis, Rutgers University.
Ross, W. D. 1930. *The Right and the Good*. Oxford: Clarendon Press.
———. 1939. *The Foundations of Ethics*. Oxford: Clarendon Press.
Russell, Bertrand. 1966. "The Elements of Ethics." In *Philosophical Essays*, 13-59. New York: Simon and Schuster.
Schroeder, Mark. 2007. "Teleology, Agent-Relative Value, and 'Good.'" *Ethics 117*: 265-95.
———. Forthcoming. "Getting Perspective on Objective Reasons." *Ethics*.
Smith, Holly M. 2010. "Subjective Rightness." *Social Philosophy and Policy* 27: 64-110.
Thomson, Judith Jarvis. 1986. "Imposing Risks." In *Rights, Restitution, and Risks*, edited by William Parent, 173-91. Cambridge, MA: Harvard University Press.
Wedgewood, Ralph. 2007. *The Nature of Normativity*. Oxford: Clarendon Press.
———. 2016. "Objective and Subjective 'Ought.'" In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman, 142-68. Oxford: Oxford University Press.
Zimmerman, Michael J. 1996. *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.
———. 2006. "Is Moral Obligation Objective or Subjective?" *Utilitas* 18: 329-61.
———. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.
———. 2014. *Ignorance and Moral Obligation*. Oxford: Oxford University Press.