

Getting Perspective on Objective Reasons*

Mark Schroeder

This article considers two important problems for the idea that what we ought to do is determined by the balance of competing reasons. The problems are distinct, but the object of the article is to explore how they admit of a single solution. It is a consequence of this solution that objective reasons—facts that count in favor—are in an important sense less objective than they have consistently been assumed to be. This raises but does not answer the question as to what evidence we ever had that reasons are as objective as has been assumed.

I. INTRODUCTION

According to a tradition descending from W. D. Ross, what we ought to do is determined by the balance of competing factors—reasons.¹ This tradition has much to be said in its favor, and it has become orthodox enough that Parfit can write that he is interested in what we ought “in the sense of most reason” to do. The idea that what we ought to do is determined by the balance of reasons is in fact so widely assumed that many philosophers, including Derek Parfit, T. M. Scanlon, and Jonathan Dancy, actually go further and claim that all other normative properties and relations

* Special thanks to audiences at the Ranch Metaphysics Workshop in January 2016 and the Lingnan Expressivism Conference in June 2016, as well as to Alex Worsnip, Daniel Wodak, Laurie Paul, Derek Baker, Shyam Nair, Ralph Wedgwood, Jake Ross, Daniel Wodak, Benjamin Kiesewetter, Nate Charlow, Janice Dowell, Karen Lewis, Ned Hall, Nina Emery, Mark van Roojen, Jamie Dreier, Paolo Santorio, Jennifer Carr, Rachael Briggs, Tom Donaldson, Trenton Merricks, Jack Spencer, Benj Hellie, and two anonymous readers and three associate editors for *Ethics*.

1. W. D. Ross, *The Right and the Good* (Cambridge: Cambridge University Press, 1930).

Ethics 128 (January 2018): 289–319

© 2018 by The University of Chicago. All rights reserved. 0014-1704/2018/12802-0001\$10.00

are determined by the balance of reasons.² This is the thesis that reasons come first. But the thesis that *ought* is analyzable in terms of reasons is subject to a number of important challenges. In this article I will be interested in two, apparently unrelated such challenges, to which I will introduce and entertain a single answer. The two birds can be killed with one stone—if we choose the right stone. If the doubly fatal answer is right, then the thesis that objective reasons are facts that count in favor has been deeply misunderstood.

The structure of the article will be simple. In Section II, I clarify the idea that *ought* is analyzable in terms of reasons, as I understand it, and distinguish between the *ought* of advisability and the *ought* of rationality, and correspondingly between objective and subjective reasons. Then in Section III, I introduce two important obstacles to this thesis, so understood—one obstacle inspired by Nico Kolodny and John MacFarlane that challenges the relationship between *ought* and reasons, and one from Ralph Wedgwood that challenges the relationship between objective and subjective reasons.³ In Section IV, I lay out a response to the second obstacle and show how it relies on *expressivism* about epistemic expressions such as ‘might’, ‘must’, and ‘likely’. In Section V, I show that the very same features of this view that allow it to constitute a response to the second obstacle extend to the first obstacle as well, allowing us to kill two birds with one stone. And finally, I close in Section VI by exploring how my solution might be generalized and considering how we could have come to be misled by talk about “objectivity” and “facts.”

II. REASONS, OUGHT, OBJECTIVE, SUBJECTIVE

A. *Reasons and Ought*

As I have already noted, the popular idea that reasons come first, among normative properties and relations, is best understood as a generalization from its core case—that what someone ought to do, in a given situation, is determined by the balance of her reasons. Ross’s original motivations for this thought are straightforward.⁴ The history of moral philosophy is full

2. Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011); T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998); T. M. Scanlon, *Being Realistic about Reasons* (Oxford: Oxford University Press, 2014); Jonathan Dancy, *Practical Reality* (Oxford: Oxford University Press, 2000); Jonathan Dancy, *Ethics without Principles* (Oxford: Oxford University Press, 2004). I myself have also defended this view, particularly in Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press, 2007); and Mark Schroeder, “Realism and Reduction: The Quest for Robustness,” *Philosophers’ Imprint* 5 (2005): 1–18.

3. Niko Kolodny and John MacFarlane, “Ifs and Oughts,” *Journal of Philosophy* 107 (2010): 115–43; Ralph Wedgwood, “The Pitfalls of ‘Reasons,’” *Philosophical Issues* 25 (2015): 123–43.

4. Ross, *Right and the Good*.

of attempts to characterize a set of absolute, unconflicting moral principles that tell us what an agent ought to do in every possible situation. But monistic theories are subject to straightforward and systematic counterexamples, and absolutist pluralistic theories must invoke complex qualifications, in order to avoid conflicts without counterexamples.

For example, the correct theory must tell us when someone ought to keep her promise to meet a friend for lunch but build in loopholes to explain why she ought not to do so when it comes into conflict with a prior, more urgent promise, or she encounters someone dying on the way to lunch whom she can save, or she gets a call that her father is on his deathbed. Similarly, it must tell us when someone ought to tell the truth but build in loopholes to explain why she ought not to do so when it comes into conflict with saving her friend's life. Moreover, the theory must explain why it is that the more significant a promise, the more significant the obstacle needs to be in order to justify breaking it, and similarly, the more significant the truth at stake, the greater the harm to one's friend needs to be in order to justify deception.

Ross's diagnosis is that the obvious explanation is that what you ought to do is the result of conflicting moral forces—what the philosophical vernacular has come to identify as *reasons*.⁵ These reasons compete, and so that is why, when you have a better reason to do something, it takes better reasons on the opposing side to justify not doing it. The fact that you have promised someone to meet her for lunch is a reason to meet her for lunch, but the fact that you have promised your spouse to always make his chemotherapy appointments may be a better reason not to meet that person for lunch, if these appointments come into conflict. And similar points go for the dying person who needs your medical attention or your father's deathbed. The reason that it is so difficult to find a general principle that incorporates all of these cases is that they do not flow from a general principle—they flow from the interaction of a variety of principles.

Ross, of course, was not the first to offer the diagnosis that the seeming complexity of the morality of promise keeping or truth telling is the result of competing factors. Sidgwick's initial argument for utilitarianism, in part 3 of *The Methods of Ethics*, is grounded in an extended argument that the best kinds of intuitive counterexamples to the principle that you should keep your promises, or to the principle that you should tell the truth, come from cases in which doing so comes into conflict with the principle of *utility*.⁶ On Sidgwick's view, there is always some utility to be gained from telling the truth or from keeping promises—it reinforces

5. Cf. esp. Dancy, *Ethics without Principles*, chap. 1.

6. Henry Sidgwick, *The Methods of Ethics*, 7th ed. (1874; repr., Indianapolis: Hackett, 1981).

the social norms of telling the truth and of keeping promises, which have great beneficial effects, for example—but occasionally there is more utility to be gained from doing otherwise. So a proper account of these cases needs to balance gains to utility on both sides—a kind of competition, as with Ross.

But Ross argued that Sidgwick's utilitarianism faces the same kind of predictable counterexamples. Just as it is predictable that we can come up with intuitive counterexamples to the principle that we should keep promises by piling up disutility, Ross argued, it is predictable that we can come up with intuitive counterexamples to utilitarianism by focusing on cases in which it comes into conflict with other important moral principles—as in cases in which it would be deeply unjust to act in a way that is only slightly better in terms of total utility. These cases, Ross argued, are just flip sides of one another. Because reasons of beneficence are among the moral reasons, the Rossian view gets a satisfying diagnosis of both sorts of conflict.

There is much more to be said about whether Sidgwick's diagnosis can be successfully generalized in response to Ross, in the form of a kind of pluralist consequentialism that resists the turn to reasons. And even if we accept Ross's view that what we ought to do is to be explained in terms of the balance of reasons, that leaves quite unsettled whether reasons themselves are to be analyzed or explained in terms of something else—value, rationality, virtue, and fittingness being among the prominent candidates. But what I want to get out of this discussion is just that the relationship between an agent's reasons and what she ought to do is supposed to be the core case for the plausibility of the thesis of Reasons First. To the extent that it is plausible that reasons come first among the normative, full stop, that is because it is compelling, in the first instance, that what we ought to do is grounded in the balance of our reasons, and because it is plausible that this line of reasoning can be extended, *inter alia*, to values, rationality, and fittingness. The challenges that I consider in this article are challenges to this core case for the centrality of reasons in moral philosophy—the idea that they determine what we ought to do.

B. Rationality, Advisability, and Objective and Subjective Reasons

We have just seen that the core case for Reasons First is the relationship between *ought* and reasons. But it helps, in order to clarify exactly what this core case is supposed to be, to follow tradition in distinguishing between two different sorts of claim that might intelligibly be made with 'ought', and two corresponding different sorts of claim about reasons. One of these sorts of claim is sensitive to what the agent herself believes, and one is sensitive only to the facts. When agents have false or incomplete beliefs, these two dimensions of evaluation can come apart. I focus first in this section, as has been done historically, on cases involving false

beliefs; both of the obstacles that I introduce in Section III turn on cases in which beliefs are incomplete.

A paradigmatic example is that of a man who (rationally) believes that the glass that he is holding contains gin and tonic, though in truth it is gasoline.⁷ If he takes a sip, then he is acting rationally, given his information. But taking a sip is not the best thing for him to do, nor what it would be appropriate for us to advise him to do. If he does take a sip, he makes a mistake. The assessment of his action as rational is sensitive to what he believes—it evaluates his action through the lens of his beliefs, or his evidence. In contrast, the assessment of his action as inadvisable, or as a mistake, is not sensitive to what he believes, or at least not in this way. It is sensitive, instead, to the facts. Some philosophers say that there are two readings of ‘ought’ claims to go along with these two dimensions of assessment—a *subjective* ‘ought’ of rationality which tells him to take a sip, and an *objective* ‘ought’ of advisability which tells him not to. The subjective ‘ought’ depends on his beliefs, but the objective ‘ought’ depends on the facts.

The way in which these two *oughts* depend on his beliefs and on the facts, respectively, is clarified by a second distinction that can be made using the same example, between two kinds of claim about reasons. The fact that his glass contains gasoline, it is said, is a reason for him not to take a sip. The sense in which this is true is the *objective* sense of ‘reason’, and it is contrasted with the sense in which this reason is available to the agent as a possible basis for action. Though it is a reason for him, it is sometimes said, it is not a reason that he has. It is an objective reason for him, but not a subjective reason for him.⁸ Different theorists adopt a variety of views about how objective and subjective reasons are related to one another—a matter to which I return shortly.

With the distinction between objective and subjective reasons in hand, the proponent of Reasons First has an elegant account of the exact way in which the objective ‘ought’ of advisability depends on the facts and the exact sense in which the subjective ‘ought’ of rationality depends on the agent’s beliefs. This is that the objective ‘ought’ says what is favored by the balance of objective reasons, and objective reasons are facts, whereas the subjective ‘ought’ says what is favored by the balance of subjective reasons, and subjective reasons are beliefs, or at least you count as having them in virtue of what you believe. In short, the idea that *oughts* are explained by reasons comes out true along both the objective/advisability and subjective/rationality dimensions of evaluation of action.⁹

7. As in Bernard Williams, “Internal and External Reasons,” in *Rational Action*, ed. Ross Harrison (Cambridge: Cambridge University Press, 1981), 101–13.

8. Cf. Mark Schroeder, “Having Reasons,” *Philosophical Studies* 139 (2008): 57–71.

9. This view is articulated particularly clearly in Mark Schroeder, “Means–End Coherence, Stringency, and Subjective Reasons,” *Philosophical Studies* 143 (2009): 223–48.

It is important for the proponent of Reasons First, however, that these two dimensions have something to do with one another. And so most—whether they accept Reasons First or not—hold that there is indeed a close relationship between objective and subjective reasons. An agent's subjective reasons, it is said, are her objective reasons that she knows.¹⁰ Or they are the considerations she takes to be objective reasons.¹¹ Or they are the propositions she believes which, if true, would be objective reasons for her.¹² All of these views take objective reasons to be prior to and explanatory of subjective reasons, and hence make the thesis of Reasons First out to be the thesis of objective reasons first.¹³ But other intimate relationships are possible. According to a natural interpretation of Dancy's view in *Practical Reality*, for example, the general notion of a reason—a core reason, as we might put it—is prior to both objective and subjective reasons.¹⁴ Objective reasons are just true core reasons, and subjective reasons are just believed core reasons. On this view, Reasons First is the thesis that core reasons come first.¹⁵ Our second challenge will strike at the idea that

10. A special case of this thesis—that your evidence is what you know—is Williamson's E=K thesis (Timothy Williamson, *Knowledge and Its Limits* [Oxford: Oxford University Press, 2000]). John Hawthorne and Jason Stanley ("Knowledge and Action," *Journal of Philosophy* 105 [2008]: 571–90) argue for the distinct but similar thesis that you should treat something as a reason just in case you know it, but they also suggest that this may be explained by the fact that your reasons are what you know, and John Hyman and Jennifer Hornsby also prominently defend the view that something can be a reason for which you act only if you know it (John Hyman, "How Knowledge Works," *Philosophical Quarterly* 49 [1999]: 433–61; Jennifer Hornsby, "Knowledge in Action," in *Action in Context*, ed. Anton Leist [Berlin: De Gruyter], 285–302). Errol Lord defends the more permissive view that your subjective reasons are the objective reasons that you are in a position to know (Errol Lord, "Having Reasons and the Factoring Account," *Philosophical Studies* 149 [2010]: 283–96).

11. Derek Parfit contrasts normative (objective) reasons with "apparent" reasons (Derek Parfit, *On What Matters* [Oxford: Oxford University Press, 2011]). Hannah Ginsborg claims that what I am calling subjective reasons always involves "taking" some consideration to be what I am here calling an objective reason (Hannah Ginsborg, "Reasons for Belief," *Philosophy and Phenomenological Research* 72 [2006]: 286–318). In a similar vein, Niko Kolodny argues that what is rational is a matter of what appears to be supported by reasons (Niko Kolodny, "Why Be Rational?," *Mind* 114 [2005]: 509–63).

12. Schroeder, "Means–End Coherence."

13. Each of these views can also be developed in ways that take different stands on the ontology of subjective reasons. According to one version of each of these views, subjective reasons are mental states; according to another version, they are propositions or facts, and there is merely a psychological condition on when these are an agent's subjective reasons. Nothing in this article turns on these distinctions.

14. Dancy, *Practical Reality*.

15. The common-core view is probably the best view out of this group, and it is motivated by distinctive problems for each of the other views. For example, the view that subjective reasons have to be known rules out the possibility that false beliefs can rationalize action, the view that subjective reasons are things you take to be objective reasons is arguably too cognitively demanding, and the subjunctive view is subject to familiar conditional

there can be any such neat relationship between objective and subjective reasons, and hence at whether reasons could have anything helpful to explain about what we rationally ought to do.

III. TWO CHALLENGES

A. *Four Envelopes*

This article concerns two challenges to the idea that what we ought to do is determined by the balance of reasons. Our first challenge is inspired by an important set of arguments from Niko Kolodny and John MacFarlane, who have given an important argument that ‘ought’, as it is used in natural language, is sensitive in some way to the speaker’s information.¹⁶ What I show in this section is that we can use cases similar to those discussed by Kolodny and MacFarlane in order to directly challenge the idea that the ‘ought’ of advisability could possibly be determined by the balance of reasons, either objective or subjective.

Earlier, we made the distinction between objective and subjective *oughts* and that between objective and subjective reasons by considering cases where an agent has a false belief. Both Kolodny and MacFarlane’s challenge and the second challenge that I introduce in the next section come from considering cases where these come apart because an agent’s beliefs are incomplete. We can best draw out the force of their challenge by considering a simplified version of one of their more complex examples. I call it the four-envelope problem, or, more generally, for reasons that will emerge, the $n + 1$ -envelope problem (fig. 1).¹⁷

fallacy—style counterexamples (cf. Robert Shope, “The Conditional Fallacy in Contemporary Philosophy,” *Journal of Philosophy* 75 [1978]: 397–413). Note that the common-core view need not carry Dancy’s specific commitments (Dancy, *Practical Reality*), which include the idea that core reasons are states of affairs and that explanations can be nonfactive.

16. Kolodny and MacFarlane, “Ifs and Oughts.” Strictly speaking, Kolodny and MacFarlane are relativists about the information sensitivity of ‘ought’ and claim that it is sensitive to the information in the context of assessment. But since speakers will say what they believe is true relative to their own context of assessment (see John MacFarlane, *Assessment Sensitivity: Relative Truth and Its Applications* [Oxford: Oxford University Press, 2014]; and Jacob Ross and Mark Schroeder, “Reversibility or Disagreement,” *Mind* 122 [2013]: 43–84), which ‘ought’ claims it is rational for speakers to make is in general sensitive to the speaker’s information.

17. The four-envelope problem is a strict generalization of what Jacob Ross (“Acceptance and Practical Reason” [PhD diss., Rutgers University, 2006]) calls the *three-envelope problem*, and which was introduced to the literature by Donald Regan’s *miners* example (Donald Regan, *Utilitarianism and Cooperation* [Oxford: Clarendon, 1980]). Setting the case up in terms of a choice between envelopes makes the structure of the problem much more transparent and also makes it easy to control for variants on the case and to generalize. My inspiration for the four-envelope problem comes from Kolodny and MacFarlane’s hydrology case, which is a variant on Regan’s original miners case. But Kolodny and MacFarlane’s case is naturally construed as relying on building in the assumption that part of the less informed observer’s information is actually false. The four-envelope case avoids this com-




FIG. 1.—Four envelopes. Color version available as an online enhancement.

Xiao faces a choice between four envelopes. She will get to open exactly one of the envelopes and will get to keep whatever she finds inside. One of the envelopes contains \$4,000, one contains \$5,000, one contains \$6,000, and the last one is empty. Xiao knows that envelope 1 contains the \$4,000, and she knows the amounts that can be found in the other envelopes, but not which envelopes they are in. Her friend Ying knows everything that Xiao knows (including what Xiao knows), and also which envelope holds the \$5,000. And their friend Zach knows everything that Ying knows, plus the contents of every envelope. Xiao knows that Ying knows which envelope holds the \$5,000, but she doesn't have a chance to communicate with him, and both know that Zach knows the contents of all of the envelopes, but neither has a chance to communicate with him.

It is easy to imagine the following three monologues:

- Xiao: The expected value of taking envelope 1 is \$4,000, but the expected value of any of the other envelopes is only \$3,667, so I should take envelope 1.
- Ying: The expected value of taking envelope 1 is \$4,000, the expected value of taking envelope 2 is \$5,000, and the expected value of taking envelope 3 or 4 is only \$3,000, so Xiao should take envelope 2.
- Zach: The expected value of taking envelope 1 is \$4,000, the expected value of taking envelope 2 is \$5,000, the expected value of taking envelope 3 is \$6,000, and the expected value of taking envelope 4 is only \$0, so Xiao should take envelope 3.

Each of these monologues makes sense in its own right—we can imagine reasoning about Xiao's choice from any of Xiao's, Ying's, or Zach's per-

plication by adding an additional option. Aaron Bronfman and Janice Dowell discuss a somewhat better case borrowed from John MacFarlane in which an Eavesdropper's judgment is informed by greater, but still incomplete, information; this case is consistent with the less informed agent having only true information, but it doesn't distinguish all three *ought* judgments, and because it involves a horse race, which is a chancy outcome, it adds a confounding role for what sorts of probabilities are involved (Aaron Bronfman and Janice Dowell, "Contextualism about Deontic Conditionals," in *Deontic Modals*, ed. Nate Charlow and Matthew Chrisman [Oxford: Oxford University Press, 2016], 117–42; MacFarlane, *Assessment Sensitivity*). The four-envelope problem removes these complications, as well.

spective, and these are plausible claims to make, given their respective states of knowledge.

Moreover, the claim of each corresponds to how it would be appropriate for them to advise Xiao about what to do, if given the opportunity. This is clear in the case of Xiao, since clearly the rational choice for Xiao to make is to take envelope 1, and it would be strange if it were not rational for her to follow her own advice. And it is clear in the case of Zach, since Zach knows all of the relevant facts, so it makes sense for him to advise her, taking only her interests into account, to take envelope 3. But it should also be clear in the case of Ying. Given what Ying knows, it would be absurd for him to advise Xiao to take anything other than envelope 2, taking only her interests into account. So each of the monologues above is plausibly interpreted as concerning the ‘ought’ of advisability.

Yet problematically, Ying’s assertion does not make sense either in terms of Xiao’s beliefs or in terms of the totality of the facts. It cannot be understood as relative to Xiao’s beliefs, because it would be irrational for Ying to make his claim under that interpretation, given what he knows. Ying knows full well that Xiao does not know which envelope contains the \$5,000, and so in the absence of that knowledge, taking envelope 2 would be an irrational risk. And yet it cannot be understood as relative to the totality of the facts, either, since it would again be an irrational claim for Ying to make given what he knows—namely, that one of the envelopes contains \$6,000. In fact, not only does Ying know this, but he also knows that somewhere Zach, who is even better informed, is thinking about Xiao’s decision and saying to himself either “She shouldn’t take envelope 2; she should take envelope 3” or “She shouldn’t take envelope 2; she should take envelope 4.” Either way, Ying is in a position already to know that taking envelope 2 is not what Xiao ought to do in light of the totality of facts about the contents of the envelopes.

Worse, Ying’s ‘ought’ claim cannot, it seems, be understood in terms of the balance of reasons at all. It cannot be the balance of subjective reasons, because it is not sensitive to Xiao’s beliefs, and it cannot be the balance of objective reasons, because it is not sensitive to the totality of the facts. Ying’s use of ‘ought’, it seems, is sensitive to his information, but this does not track either subjective reasons or objective reasons. So Ying’s case directly challenges the idea that what someone ought to do is determined by the balance of reasons, in any sense. It appears to show that the *ought* of advisability cannot be a matter of the balance of reasons, after all.

The reason that this is a significant problem is not only that the *ought* of advisability is some normative relation, and Reasons First says that all normative relations must be explained in terms of reasons. After all, there could be some indirect way in which Ying’s judgment is ex-

plained in terms of reasons, without being directly accounted for as the balance of reasons. For example, for all that I have said here, reasons could establish a utility scale, and Ying's judgment could correlate with an expectation along that scale. What makes this problem sharp, I think, is that the *ought* of advisability was supposed to be the central case for the idea that *ought* facts just are facts about the balance of reasons—the case whose compellingness is supposed to be so great that proponents of Reasons First think that it should make us interested in seeing whether it can plausibly be generalized to other cases. Yet Ying's judgment seems to track the advice that it is appropriate for him to give.

In the four-envelope problem, there are three observers, each with a distinct set of information, and all three make contrary judgments about what some agent ought to do, among a set of four choices. The challenge is that at most two of these judgments can plausibly be interpreted as rational judgments about the balance of either objective or subjective reasons. This generalizes on the more familiar three-envelope problem, sometimes known as the “miner's puzzle,”¹⁸ in which there are two distinct sets of information, and each observer makes a distinct judgment about which of three choices the agent ought to take—for example, Vivek might face a choice among three envelopes, knowing that the first contains \$2,000, one of the others contains \$3,000, and the last one is empty, without knowing which one is empty, while Wahid, who knows where the \$3,000 is, watches.

And this structure can be generalized further. Once we see how the four-envelope problem goes, we can easily see how to construct cases in which there are arbitrarily many distinct levels of relevant information, leading different agents to make contrary judgments about what someone ought to do, where at most one can be interpreted as a judgment about what she ought to do in light of her beliefs, and at most one can be interpreted as a judgment about what she ought to do in light of the totality of the facts. In general, to set up a case with n distinct levels of information and contrary judgments by n agents, it is sufficient to set up $n + 1$ envelopes, with $\$1,000 \times [\frac{1}{2}(n - 1)n + m]$ in the m th envelope for envelopes $1 - n$, and $\$0$ in the last envelope.¹⁹ The three-envelope

18. This case is modeled after the most famous example of this form, from Donald Regan, *Utilitarianism and Cooperation* (Oxford: Oxford University Press, 1980).

19. This formula guarantees that the difference between the first envelope and the empty one ($\frac{1}{2}(n - 1)n + 1$) is larger than the total sum in the other $n - 1$ envelopes in excess of the expected value of the first envelope (this sum is the $(n - 1)$ th triangular number, so it is $\frac{1}{2}(n - 1)n$). An alternative way to easily guarantee that the expected values come out right is to add more than one extra envelope for each additional thinker. For example, for a case with n levels of information, put $\$1,000 \times m$ in the m th envelope for envelopes 1 through n , and then add $\frac{1}{2}(n - 1)n + 1$ more empty envelopes.

problem and the four-envelope problem are therefore just special cases of the $n + 1$ -envelope problem.

B. *Subjective Reasons and Fineness of Grain*

Our first obstacle to Reasons First was an obstacle to the relationship between objective reasons and the *ought* of advisability. The second obstacle, due to Ralph Wedgwood, which I have also seen presented by Jacob Ross and by Daniel Wodak, is an obstacle to the relationship between objective reasons and subjective reasons.²⁰ Like our first problem, it comes from cases in which an agent's information comes apart from the facts not by being false, but by being incomplete.

Recall that though there are many different views about the relationship between objective reasons and subjective reasons, all of them are views on which there is some important correspondence between the two. Subjective reasons are known objective reasons, or they are purported objective reasons, or they are contents of belief which would, if true, be objective reasons. Or objective and subjective reasons are merely two ways in which reasons—core reasons—manifest themselves, alternatively by being true or by being believed.²¹ On every one of these views, every subjective reason is an object of belief that is a candidate to be an objective reason. But our second obstacle is an argument that this cannot be the case.

The problem is that what it is rational for an agent to do or to believe does not depend only on what she believes—it may also depend on how confident she is in one or another proposition. For example, take the case of Bernadette, who orders a gin and tonic from the bar but is four times more confident that the glass is not petrol than that it is petrol. In the jargon that has become customary, we say that she has a *credence* of .2 that the bartender has given her a glass of petrol, instead. Bernadette doesn't believe that her glass contains gin and tonic, and she doesn't believe that it contains petrol. She is more confident that it contains gin and tonic, but her .2 credence, or degree of confidence, that it is petrol makes it rational for her to set it down without taking a sip. After all, no gin and tonic is good enough to be worth a one in five chance of drinking petrol. So if what Bernadette rationally ought to do is a matter of her subjective reasons, then having a credence of .2 that the glass contains petrol must be a way for her to have a subjective reason to set the glass down without taking a sip.

20. Ralph Wedgwood, "The Pitfalls of Reasons," *Philosophical Issues* 25 (2015): 123–43; Daniel Wodak, "An Objectivist's Guide to Subjective Reasons" (unpublished); Jacob Ross, personal conversation.

21. See, e.g., Hawthorne and Stanley, "Knowledge and Action"; Ginsborg, "Reasons for Belief"; Dancy, *Practical Reality*.

But now the problem is that if we allow all possible states of credence as possible subjective reasons, we will not have enough possible objective reasons to go around, in order for every subjective reason to correspond to a possible objective reason in any of the ways that I have canvassed. Objective reasons, after all, are facts, and there are simply not enough facts to go around. Having a credence of .2 that the glass contains petrol is not, at least on the face of it, a matter of believing anything. It is not the same attitude toward a different content (such as *that it is .2 likely that the glass contains petrol*); rather, advocates of credence tell us, it is a different attitude toward the same content—that *the glass contains petrol*.²²

It is worth noting that the problem of fineness of grain is not dispelled by allowing for facts about objective chances. Though the fact that there is a .2 chance that the glass contains petrol would also be a reason not to take a sip, were it a fact, it is hard to see why this should bear on the rationality of Bernadette's choice, given that on most natural ways of filling out the case, she is certain that this is not her situation, splitting her credence instead between there being a 100% chance that the glass contains gin and there being a 100% chance that the glass contains petrol. She is sure that the world is not chancy (at least, in this respect)—her problem is that she is uncertain which unchancy way it is.

It is also worth noting that the problem of fineness of grain does admit of revisionary solutions. For example, if we assume the thesis of *uniqueness*, according to which there is only one rationally permissible prior credence function,²³ then it follows from some further common Bayesian assumptions that all rational differences in credences must be grounded in what the agent is certain of—and hence, plausibly, in what she believes. Given this package of views, we can deny that credences themselves ever rationalize belief or action and hence that they need to be subjective reasons.

Alternatively, we can reject all of the existing views about the relationship between objective and subjective reasons but defend some more relaxed, many-to-one view about this relationship. For example, instead of saying that you have a subjective reason when you believe some-

22. I have credited this problem as a problem specifically for Reasons First to Wedgwood, but Stephen Schiffer ("Review: Interest-Relative Invariantism," *Philosophy and Phenomenological Research* 75 [2007]: 188–95) raises a more specific version of this objection to Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005). Hawthorne and Stanley develop one line of response to Schiffer that can be generalized but carries very strong commitments (Hawthorne and Stanley, "Knowledge and Action").

23. For the classic statement of uniqueness, see Roger White, "Epistemic Permissivism," *Philosophical Perspectives* 19 (2005): 445–59. For general discussion, see Matthew Kopec and Michael Titelbaum, "The Uniqueness Thesis," *Philosophy Compass* 11 (2016): 189–200.

thing which, if true, is an objective reason, we could say that you have a subjective reason when you have any positive credence in something which, if true, is an objective reason. Or instead of saying that subjective reasons are things you take to be objective reasons, we could say that subjective reasons are things you have any positive credence are objective reasons.

I'm not going to argue against any of these alternative solutions to the problem of fineness of grain in this article—though some I like better than others. Instead, what I am going to do is to illustrate the virtues of a different solution—one that turns out to be surprisingly nonrevisionary, despite a strong up-front commitment. I will show how a very simple idea makes the problem go away entirely and brings with it as an almost immediate consequence a solution to the four-envelope problem, as well. It is that fact which I contend should commend it to our interest.

Like the four-envelope problem, as I have noted, the problem of fineness of grain turns on cases in which agents' information comes apart from the facts not by being mistaken, but by being incomplete. But there is no obvious *prima facie* reason to think that these are simply two faces of the same problem. The four-envelope problem requires cases of intermediate levels of information, but the problem of fineness of grain can be motivated in ordinary cases where a single agent has only partial belief about a single matter of fact. And the problem of fineness of grain requires that it is the agent of the choice herself who lacks full information, but the four-envelope problem can be fully motivated only with judgments by agents who are merely observers to someone else's choice. So they appear to be two quite different problems.

IV. FROM REASONS TO EXPRESSIVISM

A. *Triviality for Credences*

Our second obstacle claims that there are not enough possible objects of belief to go around, in order for there to be one corresponding to every possible state of credence—credences (otherwise known as degrees of belief, or degrees of confidence) are not beliefs in a wider range of contents; they are graded beliefs in the same, narrower range of contents. And since this is so, the argument alleges, there are not enough objective reasons to correspond to every possible subjective reason. As I have already noted, there is more than one place in which this argument can be resisted. But the most general, powerful response would be to make good on the idea that there really are enough objects of belief to go around. Then we could grant the objector everything but still make room for a 1–1 correspondence between objective and subjective reasons.

The problem with this is that there is a reason why the objector holds that there are not enough possible objects of belief to go around



FIG. 2.—The proposition that p . Color version available as an online enhancement.

for each credal state to correspond to some possible object of belief. It is that there is an important result which says that there is no possible object of credence, q , such that believing it is necessary and sufficient for having a credence of n in p , where p is any object of credence and n is any value in $(0, 1)$. The easiest way to see why not is pictorially. Let p be any object of credence that is not a logical truth or logical falsehood (e.g., *that the glass contains petrol*). Then we can depict this as shown in figure 2, laid out in the logical space over which credences are defined.

If q is to be an object of credence such that belief in q is necessary and sufficient for a credence of n (e.g., .2) in p , then there are five possible cases, based on the five possible logical relationships between p and q , assuming that q is also not a logical truth or a logical falsehood (fig. 3).

But now, it is easy to see that none of these choices of q can be such that belief in it is sufficient for a nonextreme credence of n (e.g., .2) in p . This is because being certain of q surely suffices for belief in q . But for each of these choices of q , someone can be certain of q without having any nonextreme credence in p . For example, in each case imagine an agent who is fully certain of the proposition represented by the dotted circle in figure 4. In every case, such an agent is certain of q and hence believes q .




FIG. 3.—Five candidates for the proposition q : that it is n likely that p . Color version available as an online enhancement.




FIG. 4.—Ways of believing q without being n confident of p . Color version available as an online enhancement.

But in case 1, every way of being certain of q entails having a credence of 0 in p , and in case 2, every way of being certain of q entails having a credence of 1 in p . In the three other cases both kinds of counterexamples are possible, but I've drawn the ones that entail a credence of 0 in p .

Pictures make the point hard to miss, but the direct proof is even simpler. Let p and q be any objects of credence and n be any value in the interval $(0, 1)$. If q is consistent with p , then there is a way of believing q —namely, by being certain of q —which entails having a credence of 1 in p , and hence which is insufficient to have a credence of n in p . And if q is consistent with $\sim p$, then there is a way of believing q —namely, by being certain of q —which entails having a credence of 0 in p , and hence which is insufficient to have a credence of n in p . But if q is not consistent with p and q is not consistent with $\sim p$, then q is itself inconsistent. So there is no consistent object of credence q such that belief in q is sufficient for a nonextreme credence of n in p , for any object of credence p .

The proof I have just given establishes a fact about the objects of credence. It shows that there is no object of credence corresponding to the content *that it is .2 likely that the glass contains petrol*, which Bernadette believes just in case she is .2 confident that the glass contains petrol. In this respect, it belongs to a family of triviality results in the spirit of Lewis.²⁴ Moreover, it is standard to assume that the objects of credence are propositions—the same as the objects of belief. But on the assump-

24. David Lewis, "Probabilities of Conditionals and Conditional Probabilities," *Philosophical Review* 85 (1976): 297–315. For general discussion of triviality results in this family, including versions of this one, see Jeffrey Russell and John Hawthorne, "General Dynamic Triviality Theorems," *Philosophical Review* 125 (2016): 307–39; see also Nate Charlow, "Triviality for Restrictor Conditionals," *Noûs* 50 (2016): 533–64.

tion that propositions are the objects of credence, the triviality result shows that there are not enough propositions to go around. But importantly, the triviality result does not show this all by itself. It only shows that there are not enough objects of credence to go around. So to get the conclusion about propositions, we need to assume that propositions are the objects of credence. It follows that the problem of fineness of grain turns on the assumption that propositions—the objects of belief—are also the objects of credence. But there is one prominent view about epistemic expressions such as ‘likely’—expressivism—on which this assumption should be denied. By adopting this view—a kind of *expressivism*—we will therefore see how to resist the problem of fineness of grain.

B. *Expressivism to the Rescue*

According to expressivism, a privileged way to characterize the meanings of the sentences of a natural language is to say what states of mind they express.²⁵ Roughly speaking,²⁶ the idea is to characterize the meaning of each sentence, *P*, by saying what it is to believe that *P*. The payoff of this approach is that it can successfully characterize the meanings of words like ‘wrong’, ‘true’, and ‘likely’ without committing to the idea, for example, that there is anything that it is for something to be wrong, anything that it is for something to be true, or anything that it is for something to be likely.

It is sometimes assumed (as I myself did, e.g., when I wrote *Being For*) that this means that expressivism is a way of doing semantics without propositions—of doing without moral propositions, or propositions concerning truth, or propositions concerning what is likely.²⁷ But this is a mistake. Expressivists can help themselves to propositions just as much as anyone else can—there is nothing about expressivism that says that there cannot be a uniform range of entities which play the role of the objects of the attitudes and the bearers of truth and falsity, and hence are what Sally believes, Harjit hopes for, and is true when Sally believes that stealing is wrong, Harjit hopes that stealing is wrong, and it is true that stealing is wrong. Since expressivists don’t accept that there is anything that it is for stealing to be wrong, however, they will simply accept unorthodox views about the nature of propositions. For example, they will deny that in order to do their job, propositions need to determine

25. For the canonical statement, see Allan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 1990).

26. Greater finesse in this formulation is required in order to deal with slurs, conventional implicatures, presuppositions, and—saliently, in the context of this article—some treatments of epistemic expressions, such as that given by Seth Yalcin, “Epistemic Modals,” *Mind* 122 (2007): 867–914.

27. Mark Schroeder, *Being For: Evaluating the Semantic Program of Expressivism* (Oxford: Oxford University Press, 2008).

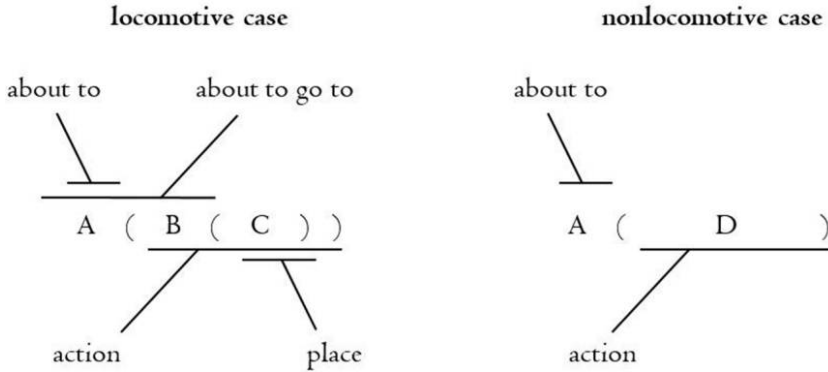


FIG. 5.—The state of being about to go somewhere.

their own truth conditions, in any interesting sense, or to be intrinsically representational.²⁸

The proper expressivist attitude toward orthodox theories of propositions is that they are based on a simple conflation that results from theorists considering a limited diet of cases.²⁹ The belief that grass is green, on this view, is like the state of being about to go to Paris. It can be carved up as a relation to a place—Paris—the relationship of being about to go somewhere. Or it can be carved up as a relation to an action—going to Paris—the relationship of being about to do it (fig. 5).

Orthodox theorists have argued that belief is a relationship to propositions, argued that paradigmatic beliefs such as the belief that grass is green can be characterized by a relation to intrinsically representational entities that determine their own truth conditions, and concluded that propositions are intrinsically representational entities that determine their own truth conditions. But this is like arguing that being about to go to Paris is a relation to an action, arguing that it is a relation to a place, and concluding that actions are or determine places. As the expressivist sees things, orthodox theorists have located propositions in the wrong place. They have taken propositions to be places (in the analogy), when really they are actions. By restricting their attention to what is merely a special case of belief, they have misidentified propositions with some-

28. So expressivists will reject orthodox assumptions about the nature of propositions that are presupposed by contemporary work on the nature of propositions such as that in, e.g., Jeff King, *The Nature and Structure of Content* (Oxford: Oxford University Press, 2007); Scott Soames, *What Is Meaning?* (Princeton, NJ: Princeton University Press, 2012); Greg Ackerman, "Representation, Truth, and the Metaphysics of Propositions" (PhD diss., University of Southern California, 2016).

29. For more on this expressivist perspective on propositions, see particularly Mark Schroeder, "Two Roles for Propositions: Cause for Divorce?," *Noûs* 47 (2013): 409–30; and the other essays collected in Mark Schroeder, *Expressing Our Attitudes* (Oxford: Oxford University Press, 2015). The diagrams come from "Two Roles for Propositions."

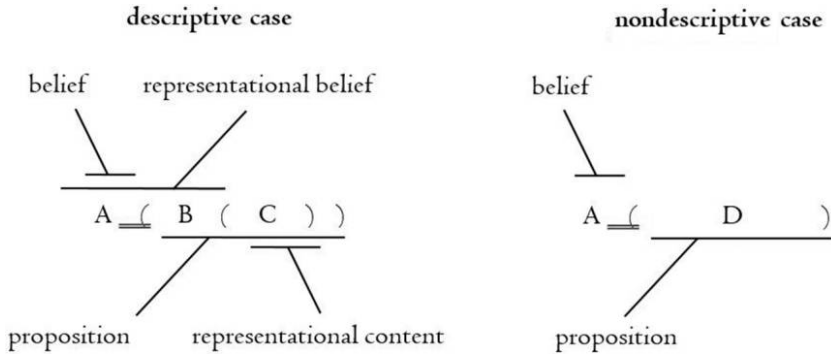


FIG. 6.—The state of believing something.

thing that is insufficiently general. Just as not all actions consist in going someplace, not all propositions consist in relations to intrinsically representational contents (fig. 6).

For the expressivist about epistemic expressions, therefore, propositions cannot be identified with the objects of credence—they are something more general. Some propositions do determine objects of credence—these are ordinary descriptive propositions. To believe such a proposition is therefore to be in both a relation to a proposition and a relation to an object of credence. But the relation that you are in to the object of credence is not belief, any more than the relation that you stand in to Paris is the relation of being about to do it. The belief relation is a more general relation, which you can stand in to any proposition. And some propositions—such as the proposition *that it is .2 likely that the glass contains petrol*—are nondescriptive. They do not correspond to any possible object of credence. But that is no obstacle to their being propositions—objects of belief—because not all objects of belief determine objects of credence.

A familiar way in which such a view can be motivated is through standard “triviality” results about conditionals. After Ernest Adams postulated that the assertability of the conditional ‘if *P*, then *Q*’ should track the speaker’s conditional credence in *Q*, conditional on *P*, and Robert Stalnaker postulated that the proposition expressed by the conditional is the one such that your credence in it matches your conditional credence in *Q*, conditional on *P*, David Lewis famously showed that there is no possible object of credence such that your credence in that content matches your conditional confidence in *Q*, conditional on *P*.³⁰ Theorists

30. Ernest Adams, “The Logic of Conditionals,” *Inquiry* 8 (1965): 166–97; Ernest Adams, *The Logic of Conditionals: An Application of Probability to Deductive Logic* (Dordrecht:

about conditionals who follow Adams's idea that your confidence that "if P , then Q " should match your conditional confidence in Q conditional on P , including Dorothy Edgington and Jonathan Bennett, have typically denied, for this reason, that conditionals express propositions.³¹ But they do not need to say this; they can conclude, instead, that propositions are not the objects of credence.

And we have already noted that the result in Section IV.A is just a special case of a triviality result—one for credences, rather than for conditional belief. It shows that there is no object of credence q such that belief in q can be identified with a credence of n in p . We could conclude from this that there is no proposition q such that belief in q can be identified with a credence of n in p . But we should not. Expressivists can deny this, because their view is naturally construed as one on which we take a more general perspective on the nature of propositions.³²

Because this move is so high-level, it is worth spelling out how this changes the triviality arguments. If sentences containing epistemic expressions such as 'probably' and 'it is .2 likely that' express propositions that do not correspond to objects of credence, then what goes wrong in triviality arguments, which generally do not (as mine does not) mention the word 'proposition'? The answer is that triviality arguments all work by assuming that sentences containing the target expression ('if . . . then', 'probably', 'it is .2 likely that') are subject to the usual rules of the probability calculus—for example, that probabilities (corresponding to credences) over these are well defined, that they can be conditionalized on, and that these probabilities obey the normal laws of probability. But that is precisely what we deny, by saying that the propositions expressed by these sentences do not correspond to objects of credence. Either there is no such thing as being more or less confident in them, or degrees of confidence in them do not behave like probabilities.³³

Reidel, 1975); Robert Stalnaker, "Probability and Conditionals," *Philosophy of Science* 37 (1970): 64–80; Lewis, "Probabilities of Conditionals."

31. Dorothy Edgington, "Do Conditionals Have Truth Conditions?," *Critica* 18 (1986): 3–30; Dorothy Edgington, "On Conditionals," *Mind* 104 (1995): 235–329; Jonathan Bennett, *A Philosophical Guide to Conditionals* (Oxford: Oxford University Press, 2003).

32. Not all expressivists about epistemic expressions describe their views in this way, but those who do not, such as Yalcin ("Epistemic Modals"), still have objects in their theories that play all of the theoretical roles of propositions but do not determine sets of possible worlds—they just leave the word 'proposition' to their orthodox opponents. The right way to think about these theorists is as believing that propositions are more general than objects of credence but as having misgivings about using the word 'proposition' to describe this, given that it has been taken over by orthodoxy.

33. For example, the friend of Adams's thesis could hold that confidence in a conditional behaves like a conditional probability, rather than like a probability. Or if believing that it is .2 likely that p is just having a .2 credence in p , then there may be no such thing as being more or less confident that it is .2 likely that p .

V. TWO BIRDS WITH ONE STONE

A. *The Four-Envelope Problem*

So far, we have seen the main ingredient for how an expressivist about epistemic expressions can answer the problem of fineness of grain. On this view, there are enough possible objective reasons to correspond to every possible subjective reason, because objective reasons are facts, facts are true propositions, and every state of credence is a matter of believing some proposition. So when Bernadette has a subjective reason not to take a sip in virtue of having a credence of .2 that the glass contains petrol, there is some proposition that she believes—namely, *that it is .2 likely that the glass contains petrol*—which, if true, is an objective reason not to take a sip. What I will now show is that this resolution of the problem of fineness of grain has the virtue that it also answers our first obstacle—the four-envelope problem.³⁴

The reason that the expressivist's response to the problem of fineness of grain extends to the four-envelope problem has to do with what it requires the expressivist to say about objective reason judgments. The point of the problem of fineness of grain is not just to locate propositions corresponding to every possible state of credence; it is to locate propositions which could themselves be candidate objective reasons, if they are true. But there is no perspective-independent fact about the world which determines whether the proposition *that it is .2 likely that the glass contains petrol* is true. There are only facts about who believes it to be true and who does not—it is believed to be true by people whose credence that the glass contains petrol is .2, and not otherwise.

Similarly, if this proposition is an objective reason, if true, then there is no perspective-independent fact about the world which determines whether it is an objective reason or not. There are only facts about who is committed to believing that it is an objective reason. So it follows, on this view, that which judgments it makes sense for someone to make about objective reasons is sensitive to her state of information. But the four-envelope problem is precisely the problem that which judgments it makes sense for someone to make about the 'ought' of advisability is sensitive to her state of information. But given the expressivist response to the problem of fineness of grain, this is precisely what we should ex-

34. It should go without saying, but perhaps it does not, that epistemic expressivism and norm expressivism are independent of one another. The two-bird solution being explored in this article is expressivist about epistemic expressions like 'might', 'must', 'if . . . then', and '.2 likely' and therefore expressivist about 'ought' and 'objective reason' because they fall under this class. This has nothing to do with norm expressivism, the more familiar form of expressivism in ethics, except that someone who endorses epistemic expressivism has limited resources for objecting to norm expressivism on grounds of the Frege-Geach problem. See Mark Schroeder, "Attitudes and Epistemics," in Schroeder, *Expressing Our Attitudes*.

pect, if what someone ought—in the sense of advisability—to do is determined by the balance of her objective reasons.

We can see the same point working backward from the structure of the four-envelope problem. The crux of the four-envelope problem is that Ying's 'ought' judgment depends on his own knowledge or belief state, and not on either Xiao's beliefs or the totality of the facts. This is what seems to prevent it from being grounded in reasons. For it to be grounded in reasons, Xiao's reasons would have to depend on Ying's beliefs! But this seems absurd. Surely, Xiao's reasons are exhausted by her subjective reasons, which depend only on Xiao's beliefs (and not on Ying's), and her objective reasons, which depend only on the facts (and not on Ying's beliefs). So it seems that Ying's 'ought' judgment cannot be linked to reasons. So, at any rate, this is how we reasoned in setting up the problem.

But notice that if we are epistemic expressivists, then among the propositions that we acknowledge is the proposition that the expected value of Xiao's taking envelope 3 is \$3,000. And so if facts are just true propositions, then if this proposition is true, it is a fact, and hence a candidate to be an objective reason (which, after all, are just facts, as is agreed on all sides). Of course, given my description of the case, we all know that it is not true—since we know that the \$6,000 is in envelope 3, we know that the expected value of Xiao's taking envelope 3 is \$6,000, rather than \$3,000. But Ying doesn't know this. Ying believes that the expected value of Xiao's taking envelope 3 is \$3,000, because he believes that it is .5 likely that it contains \$6,000 and it is .5 likely that it is empty. (And according to the epistemic expressivist, he counts as believing these things because his credence that envelope 3 contains \$6,000 is .5, as is his credence that it is empty.) So Ying believes that it is a fact that the expected value of Xiao's taking envelope 3 is \$3,000. So it does make sense for him to believe that there is an objective reason for Xiao to take envelope 3, though not as good as the objective reasons for Xiao to take envelope 1 or 2. Since he thinks the same about envelope 4, it therefore makes perfect sense for him to believe that taking envelope 2 is what is supported by the balance of Xiao's objective reasons—or, in short, that she objectively ought to take envelope 2.

In setting up the four-envelope problem, we considered the three monologues that characterize reasonable things each of Xiao, Ying, and Zach could be thinking about Xiao's choice. We can now expand on these monologues, to see how these judgments are really judgments about the balance of objective reasons:

- Xiao: The expected value of envelope 1 is \$4,000. That is an objective reason to take envelope 1.
The expected value of envelope 2 is \$3,667. That is an objective reason to take envelope 2.

The expected value of envelope 3 is \$3,667. That is an objective reason to take envelope 3.

The expected value of envelope 4 is \$3,667. That is an objective reason to take envelope 4.

The objective reason to take envelope 1 is better than the objective reasons to take envelope 2, 3, or 4.

So I should (objectively) take envelope 1.

Ying: The expected value of envelope 1 is \$4,000. That is an objective reason to take envelope 1.

The expected value of envelope 2 is \$5,000. That is an objective reason to take envelope 2.

The expected value of envelope 3 is \$3,000. That is an objective reason to take envelope 3.

The expected value of envelope 4 is \$3,000. That is an objective reason to take envelope 4.

The objective reason to take envelope 2 is better than the objective reasons to take envelope 1, 3, or 4.

So Xiao should (objectively) take envelope 2.

Zach: The expected value of envelope 1 is \$4,000. That is an objective reason to take envelope 1.

The expected value of envelope 2 is \$5,000. That is an objective reason to take envelope 2.

The expected value of envelope 3 is \$6,000. That is an objective reason to take envelope 3.

The expected value of envelope 4 is \$0. That is not an objective reason to take envelope 4.

The objective reason to take envelope 3 is better than the objective reasons to take envelope 1, 2, or 4.

So Xiao should (objectively) take envelope 3.

Before moving on, it is important to bring out one important feature of this solution. I have been treating all three of Xiao's, Ying's, and Zach's claims as claims about what Xiao objectively ought to do, in the sense of what is determined by the balance of all of her objective reasons. This choice is not arbitrary. In order to solve the problem of fineness of grain, it was not enough just to accept expressivism about epistemic expressions; we also had to accept that some epistemic propositions are, if true, objective reasons. It is an immediate consequence of this that what it is to judge that the objective reasons in some situation are such-and-such must be in part a matter of your credal state, or of the evidence that you have available to you. So except in very special cases, generalizations about the totality of objective reasons will be sensitive to the speaker's evidence. But according to the view that *oughts* are determined by reasons, objective *ought* judgments just are generalizations about the totality of objective reasons. So we should expect them to behave like this.

B. All of the Facts

This expressivist solution to the four-envelope problem claims that Ying's judgment is an objective 'ought' judgment after all—a judgment about what Xiao ought to do in light of all of the objective reasons. It is not, of course, a judgment about what she ought to do in light of all of the facts, because there is a fact about which envelope holds the \$6,000, and Ying knows that there is such a fact. What Ying cannot accept, therefore, is the following line of reasoning:

1. Either the \$6,000 is in envelope 3, or the \$6,000 is in envelope 4.
2. If the \$6,000 is in envelope 3, then that is an objective reason for Xiao to take envelope 3 that outweighs any objective reason for her to take envelope 2.
3. If there is an objective reason for Xiao to take envelope 3 that outweighs any objective reason for her to take envelope 2, then the balance of objective reasons does not favor Xiao taking envelope 2.
4. If the \$6,000 is in envelope 4, then that is an objective reason for Xiao to take envelope 4 that outweighs any objective reason for her to take envelope 2.
5. If there is an objective reason for Xiao to take envelope 4 that outweighs any objective reason for her to take envelope 2, then the balance of objective reasons does not favor Xiao taking envelope 2.
6. So reasoning by cases, either way the balance of objective reasons does not favor Xiao taking envelope 2.
7. So it is not the case that Xiao objectively ought to take envelope 2.

This reasoning feels good. It feels like Ying should be able to reason in this way. In fact, this is the reasoning that we relied on in setting up the four-envelope problem, in the first place—it was the reason it didn't seem like Ying's judgment could be a (rational) judgment about what Xiao objectively ought to do. If you are tempted by this line of thought, then you are thinking that the expressivist solution to the four-envelope problem is absurd, since that solution turns on the view that what Ying is judging is that the balance of all objective reasons does favor Xiao taking envelope 2.

But in the context of expressivism we cannot trust this reasoning. For example, Ying would make a mistake to reason as follows:

1. Either the \$6,000 is in envelope 3, or the \$6,000 is in envelope 4.
2. If the \$6,000 is in envelope 3, then Zach knows that it is, and hence knows that it is not the case that the \$6,000 might be in envelope 4.

3. If Zach knows that it is not the case that the \$6,000 might be in envelope 4, then he knows that it is not the case that the \$6,000 might be in envelope 3 and might be in envelope 4.
4. If the \$6,000 is in envelope 4, then Zach knows that it is, and hence knows that it is not the case that the \$6,000 might be in envelope 3.
5. If Zach knows that it is not the case that the \$6,000 might be in envelope 3, then he knows that it is not the case that the \$6,000 might be in envelope 3 and might be in envelope 4.
6. So reasoning by cases, either way Zach knows that it is not the case that the \$6,000 might be in envelope 3 and it might be in envelope 4.
7. So it is not the case that the \$6,000 might be in envelope 3 and it might be in envelope 4.

Obviously something goes wrong with this piece of reasoning. Given his evidence, Ying should think that the \$6,000 might be in envelope 3 and it might be in envelope 4, and so he should definitely not conclude by such reasoning that this is not the case.

Importantly, as Jacob Ross and I have shown, every kind of information-sensitive or *epistemic* expression can be used in at least some sentences such that it makes perfect sense for a rational thinker to assert or accept these sentences even though she knows that someone who knows everything that she does plus more—even her own future, more informed self—denies this very sentence.³⁵ We call these cases of *rational reversibility*, and they can be constructed for every arguably information-sensitive expression—‘might’, ‘must’, ‘probably’, ‘it is .2 likely that’, and the indicative conditional, among others. And every case of rational reversibility is a case in which reasoning by cases, like each of the kinds that Ying might go through, above, fails.

Importantly, therefore, this is a *general* feature of epistemic expressions. Once we recognize that ‘objective reason’ belongs to the class of epistemic expressions, the fact that we cannot reason with them in this way should not be surprising—it is a general consequence of the fact that it is never safe to reason with epistemic expressions in this way. As soon as we grant that there can be objective reasons like the fact that it is .2 likely that the glass contains petrol, and that this is an objective reason only if it is .2 likely that the glass contains petrol, it follows that ‘objective reason’ falls under the class of expressions with which we cannot safely reason in this way:

1. Either the glass contains petrol, or it doesn’t.
2. If the glass contains petrol, then it must contain petrol.

35. Ross and Schroeder, “Reversibility or Disagreement.”

3. If the glass must contain petrol, then it is not .2 likely to contain petrol, and so it is not true that the fact that it is .2 likely to contain petrol is an objective reason not to take a sip.
4. If the glass doesn't contain petrol, then it must not contain petrol.
5. If the glass must not contain petrol, then it is not .2 likely to contain petrol, and so it is not true that the fact that it is .2 likely to contain petrol is an objective reason not to take a sip.
6. So reasoning by cases, either way it is not true that the fact that it is .2 likely to contain petrol is an objective reason not to take a sip.

So we know that in general, reasoning by cases must fail with claims about objective reasons—this follows as soon as we take the expressivist solution to the problem of fineness of grain. The only way, therefore, that Ying's reasoning about the balance of all objective reasons could be good is if there is something very special about his context or about the exact conclusion that he is drawing that makes his particular use of reasoning by cases acceptable, even though reasoning by cases is in general bad when the conclusion concerns objective reasons. The picture that would be needed in order to validate Ying's problematic reasoning, therefore, is given by figure 7.

This is not an incoherent picture. For example, it is possible to hold that though considerations like the fact that it is .2 likely that the glass contains petrol can be objective reasons, these objective reasons never matter for the balance of all objective reasons, because they are always swamped by the objective reasons that come from nonepistemic facts—facts like the fact that the glass does in fact contain petrol, or the fact that it does not contain petrol. Someone who endorses this swamping view could embrace the expressivist solution to the problem of fineness of grain but resist the expressivist solution to the four-envelope problem, by offering an explanation of why Ying can apply reasoning by cases to




FIG. 7.—Reasoning by cases.

conclude that it is not the case that Xiao objectively ought to take envelope 2.

What this shows is that expressivism about epistemic expressions—the stone that kills the problem of fineness of grain—does not automatically kill the four-envelope problem. But it does not show that it is not the right kind of stone to kill the four-envelope problem. On the contrary, once we are expressivists about epistemic expressions and agree that among the objective reasons are facts like that it is .2 likely that the glass contains petrol—the two assumptions required to solve the problem of fineness of grain—we get a solution to the four-envelope problem unless we adopt special assumptions like the one above in order to try to resist this solution. And so my question is, why would we want to resist this solution?

Given Kolodny and MacFarlane's arguments that the same pattern of reasoning fails for the 'ought' of advisability, we should have expected that it would have to fail for 'objective reason' judgments as well, if there is any intimate connection between objective reasons and the 'ought' of advisability.³⁶ So whatever feels surprising about the result that Ying cannot reason his way by cases to the conclusion that Xiao ought not to take envelope 2 is arguably just a product of the fact that everyone has been assuming all along that objective reason judgments are not information sensitive. On the expressivist view being considered here, they are.

VI. TIDYING UP

A. *Objective Reasons and Objective Ought: Less Objective Than Most Have Thought*

The solutions that I have been introducing in this article to the four-envelope problem and the problem of fineness of grain entail that both objective reasons and objective 'ought' judgments are along a certain dimension less objective than has been nearly universally assumed throughout moral philosophy. When cases like Williams's classic gin and tonic case have been used to distinguish the 'ought' of advisability from the 'ought' of rationality, it has seemed just as obvious that the objective 'ought' judgments depend simply on the totality of the facts as that the subjective 'ought' judgments depend simply on the information available to the agent of the choice. What we have seen in this article is that, at least according to the view that I have been exploring here, though this thought is in a sense correct, it has also been misunderstood.

On the expressivist solution being explored in this article, objective 'ought' judgments do depend on the totality of the facts—but with two important qualifications. The first is that there are more facts on which

36. Kolodny and MacFarlane, "Ifs and Oughts."

objective 'ought' judgments can depend than has previously been assumed, because there are possible facts such as *that it is .2 likely that the glass contains petrol*. This is the upshot of the expressivist's answer to the problem of fineness of grain. The second is that not all facts matter—only the facts that are reasons. This is the upshot of the expressivist's answer to the reasoning by cases that helps to drive the four-envelope problem. The expressivist does not deny that there is a fact about which envelope contains the \$6,000, nor that if it is in envelope 3, that is a conclusive objective reason to take envelope 3, nor that if it is in 4, that is a conclusive objective reason to take envelope 4. The expressivist who shares Ying's information, however, denies that it follows from these three claims that either there is a conclusive objective reason to take envelope 3 or there is a conclusive objective reason to take envelope 4. The facts are still there, but that is not enough for them to provide reasons—only conditional reasons.

Previous theorists have missed these qualifications, but they have otherwise latched onto central claims that are importantly true: first, that what an agent objectively ought to do is determined by the balance of the totality of objective reasons, and second, that objective reasons are facts which count in favor. Both of these claims are retained by the expressivist view. So though it can sound surprising or unintuitive, the expressivist view is actually a very conservative view about objective reasons and the relationship between objective reasons and the objective 'ought'. It even accepts both of these very claims that (according to it) other theorists have misunderstood. All that it gives up is a common interpretation of the upshot of these central truths about objective reasons and the objective 'ought'—an upshot that is mediated by a background theory about what sorts of things can be facts.

And this, I think, is exactly what we should hope for. For despite the fact that the expressivist's account of the objective 'ought' makes it out to be somewhat less objective than orthodoxy would have led us to expect, the expressivist's objective 'ought' actually has just the right amount of objectivity to be helpful for deliberation and advice. Taking the case of deliberation first, if 'ought' were more objective, then Xiao should conclude that objectively she should not take envelope 1. This is what she knows is believed by people who are better informed than she is. But if this is true, then the objective 'ought' is not relevant for deliberation. Similarly, if 'ought' were more objective, then Ying should conclude that objectively Xiao should not take envelope 2. That is what he knows is believed by someone who is better informed than he is. But if this is true, then the objective 'ought' is not relevant for how Ying should advise Xiao.

If the objective 'ought' is not useful for either deliberation or advice, then it lacks the central features that might make it practical. In-

deed, the first half of this point—the point about deliberative irrelevance—is the exact argument given by many philosophers that the objective ‘ought’ is normatively unimportant!³⁷ So rather than seeing the expressivist solution in this article as one on which we lose something important about the objective ‘ought’ because it turns out to be less objective than we thought, we should actually see it as one on which we get something back about the importance of this central normative notion.

After introducing the distinction between objective and subjective ‘ought’ judgments in the *Methods of Ethics*, Sidgwick made the highly natural claim that these never come apart from the first-person deliberative perspective. Sidgwick’s thought is natural, because if we restrict our attention to cases in which what an agent subjectively ought to do comes apart from what she objectively ought to do only because one of her beliefs is false, the agent herself will never rationally and reflectively believe that she is currently in such a situation.

But once we recognize that what an agent subjectively ought to do can come apart from what she objectively ought to do, in addition, because her beliefs are incomplete, his thought is hard to maintain. In three-envelope cases, it is rational for an agent to take the choice that she knows is only second best, simply because the risk of the worst outcome is not worth the gain of the best outcome. In these cases, reasoning by cases leads us to think that the agent herself can know that she objectively ought not to take the first envelope, even though subjectively she ought to. But the expressivist view being explored in this article lets us get Sidgwick’s idea back. On the expressivist view, this reasoning by cases fails, and hence no rational and reflective agent ever faces a conflict between her beliefs about what she objectively and subjectively ought to do. So Sidgwick turns out to have been right after all.

If you are like most readers, at this point you are probably still deeply skeptical about whether it makes sense to think that Ying’s judgment concerns the objective ‘ought’. If so, then you are probably being led along by some tacit theory. We originally used cases like Williams’s gin and tonic case to distinguish two things that are worth saying about what someone “ought” to do—one of which depends on her beliefs, and one of which depends on the facts of the case. That much is clear and I think pre-theoretical. Theorists have jumped to a conclusion about how the latter sort of judgment depends on the facts of the case, starting from a naive assumption about what sort of things can be facts. That was a natural jump, and it could be right, for all that I have said in this article—the expressivist stone, as attractive as its bird-killing benefits might be, could be false. But what makes you so sure?

37. Ross, “Acceptance and Practical Reason”; John Gibbons, “You Gotta Do What You Gotta Do,” *Noûs* 43 (2009): 157–77.

It makes sense to be sure about what is built into the meaning of a stipulative use of ‘objectively ought’, and for many philosophers by this point, the theory is old enough to have been replaced with stipulation. But it does not make sense to be sure about whether such a stipulative use tracks anything of pretheoretical interest. The expressivist view does not claim that Ying’s judgment is a judgment about what Xiao “objectively ought” to do in the stipulative sense of what depends only on information-invariant facts that is tacit in much contemporary theory. But it does claim that Ying’s judgment is a judgment about what Xiao objectively ought to do in the sense that was in play in the ordinary ‘ought’ judgments in cases like Williams’s that were originally used to introduce the ‘objective’/‘subjective’ distinction in the first place—that is, the sense of what is supported by the balance of all of Xiao’s objective reasons. And it has an explanation of how other theorists could have started by theorizing about this sense of ‘ought’ and been diverted by way of a natural but false assumption about what sort of thing can be reasons. To evaluate it, we have to get past the thought that it is not how we have been thinking of things and engage with it on its own terms.

B. Postscript: Doing Away with the Ladder

In this article, I have been showing how to kill two birds with one stone and thereby give elegant and surprisingly conservative answers to two important and underappreciated challenges to the much-discussed idea that what someone ought to do is determined by the balance of reasons. The problem, some will say—and I think rightly—is that these solutions do not come without commitments. On the contrary, they flow from the thesis of expressivism about epistemic vocabulary. And expressivism about epistemic vocabulary is itself a controversial thesis. So it is worth pausing at least briefly to note where expressivism has done the heavy lifting and to consider whether this element of the solutions explored in this article might ultimately be replaceable within the framework of an alternative view about epistemic expressions, such as relativism or contextualism.

One reason why you might be initially optimistic that relativism or contextualism about epistemic expressions could do similar work to expressivism, in the context of the problems considered in this article, is that all three views in this triad will have (competing) explanations for why reasoning by cases generally fails with epistemic vocabulary. So if ‘objective reason’ turns out to be an epistemic expression on a par with ‘objectively ought’, then these will pattern together on any of these general theories of epistemic expressions, and we can explain why the reasoning by cases will fail, and hence why we can recover the idea that Xiao’s and Ying’s judgments belong to the same category as Zach’s—all three are filtered through the agent’s own information.

Expressivism, therefore, is not required in order to get this far with the four-envelope problem. Where expressivism has come in, in this article, is in giving an explanation of why we should expect 'objective reason' to be an epistemic expression, and by giving us a candidate for what the objective reasons could be. Even without the help of expressivism, we could describe a semantics for 'objective reason' on which 'there is an objective reason for Xiao to take envelope 2 that is better than an objective reason for her to take envelope 3 or 4' is something that it is rational for Ying to accept given his knowledge. But it is hard to say what Ying should think that this reason is. Similarly, it is not clear why objective 'reason' judgments are information sensitive in this way.

Expressivism gives these two questions the same answer. It is common ground that an objective reason must be a truth. So since the reason for Xiao to take envelope 2 is that its expected value is \$5,000, and this claim is an epistemic one, the claim that it is an objective reason—and hence even the claim that there is an objective reason—must consequently be epistemic, as well. And these two answers march hand in hand with the expressivist's answer to the problem of fineness of grain, since, according to the expressivist, Bernadette has a credence of .2 that the glass contains petrol just in case she believes that it is .2 likely that the glass contains petrol—which is the very consideration that the expressivist claims is a reason for her not to take a sip.

Contextualist and relativist views are going to struggle with making good on each of these pieces of my two-birds-with-one-stone solution. According to both contextualist and relativist views, for example, there is an ordinary proposition expressed by 'it is .2 likely that the glass contains petrol' in any given context of utterance, but it is not one that the agent believes just in case she has a credence of .2 that the glass contains petrol; rather, it will be one that the agent will believe just in case she believes that she has such a credence, or that such a credence would be appropriate in her situation. Since these two can come apart, the contextualist and relativist attempts to implement this sort of solution will fail to make good on exactly the ways in which nonextreme credences are held to rationalize action by the proponent of the objection.

And likewise, contextualist and relativist views will struggle with identifying the reason for Xiao to take envelope 2, which Ying accepts. According to the contextualist, for example, the proposition expressed by 'the expected value of her taking envelope 2 is \$5,000' will vary from context to context, but it will in general be either a proposition about the expected value conditional on some particular body of information which happens to be Ying's or a proposition about the expected value conditional on Ying's information.³⁸ But neither of these propositions is itself informa-

38. Or possibly: some body of information that includes Ying's, or the like.

tion sensitive. And this creates a problem, because since these propositions are supposed to be, if true, objective reasons, that creates an overabundance of objective reasons. After all, for the contextualist, all three of Xiao, Ying, and Zach might speak truly when they say, respectively, ‘the expected value of envelope 3 is \$3,667’, ‘the expected value of envelope 3 is \$3,000’, and ‘the expected value of envelope 3 is \$6,000’. So if all of those propositions are, if true, objective reasons, then there are a lot of objective reasons for Xiao to take envelope 3, and Ying is not in such a great position, after all, to be rationally confident that the balance of objective reasons supports taking envelope 2.

I don’t want to argue here that neither contextualism nor relativism can rescue the solution that I have been offering, with a suitable song and dance. I just mean to be illustrating that doing so is nontrivial—expressivism actually does real work in the response as I have developed it. But the fact that it does real work does not mean that it absolutely cannot be mimicked. Perhaps it can, and speaking for myself, I hope that it can, since I think that the evidence for contextualism is, on the whole, better than the evidence for expressivism. But even if expressivism about epistemic expressions is not the truth, it shows us exactly what contextualism or relativism would need to find a way to mimic, in order to kill both birds with one stone and save the day for the Rossian idea that what we ought to do is determined by the balance of reasons, which is so central to the pervasive idea that reasons have a central place in moral philosophy.