

TRUTH AND CORRECT BELIEF

Allan Gibbard
University of Michigan

For belief, correctness is truth. Correct belief is true belief. My belief that snow is white is correct just in case the belief is true, just in case snow is white. Correctness, now, seems normative. More precisely, as we should put it, the concept of correctness seems to be a normative concept—and that raises a puzzle: Hume worried about the transition from *is* to *ought*, and the lesson that many have drawn is that from purely non-normative premises alone, no normative conclusion can follow. Non-normative facts do support normative conclusions, to be sure, as the fact of rain can support the conclusion that I ought to take an umbrella. The support, though, is not a pure matter of logic, and it isn't analytic. We can't deduce with sheer logic and an understanding of meanings that rain supports taking an umbrella; we rely on the further implicit premises that in rain the umbrella prevents misery and that misery is, other things equal, to be avoided. Yet from the truth of "Snow is white" follows the correctness of a belief that snow is white.

Perhaps, one might think, this follows because the notion of truth itself is a normative notion. That, however, won't resolve the puzzle. To say that it's true that snow is white amounts just to saying that snow is white. From the fact that snow is white, then, it follows that the belief that snow is white is correct. That snow is white is a non-normative claim if any claim is non-normative. It seems, then, that in this case, a normative conclusion follows from a non-normative premise alone. It's not just that the whiteness of snow supports the conclusion, as rain supports taking an umbrella; the relation is more intimate than that. That snow is white entails analytically, it seems, a normative conclusion.

It would of course be good to explain this puzzle away, to say why there is really no puzzle. And explaining the puzzle away might deliver valuable lessons. The seeming puzzle shows, perhaps, that the whole normative/non-normative

gap is misconceived, that what we think non-normative mixes without boundary with what we think most normative. That would be one way to explain the puzzle away. I want to experiment, though, with maintaining a normative/non-normative distinction. Normative findings are matters of what to do and the like, and questions of what to do seem different from questions of how things are. In the case of belief, to be sure, normative and non-normative questions will be intimately related: what to believe ties in intimately with how things are. Whether snow is white settles, in a sense, whether to believe that snow is white. I'll be exploring, though, whether we can maintain a normative/non-normative distinction even for correctness of belief. Can we distinguish questions of what to believe from questions of how things are, and yet explain their analytic equivalence?

Whether a belief is correct is a different question from whether it is advantageous or whether it is desirable. In the stock case of the spouse who may be having an affair, the comforting belief may not be the correct one—and if comfort isn't the first desideratum, that's not a matter of the logic of belief. What, then, does it mean to say that a belief is correct?

We might try distinguishing what it's *correct* to believe from what one *ought* to believe. No *ought*, perhaps, follows from a non-normative *is*, but a normative *correct* can follow.¹ Now the term 'ought' can have many meanings, but I'll try to elucidate one meaning that ties in intimately with correctness. What one ought to do, all things considered, we might also speak of as "the thing to do" or what it "makes most sense" to do. It is what there's most reason to do, all things considered. A like sense of 'ought' applies to preferences and to beliefs: What one "ought" in this sense to prefer is what's preferable. What's preferable need not be what it's preferable to prefer: for the prisoner without hope of escape, for instance, preferring freedom may only bring frustration, in which case it may be preferable to prefer confinement. Still, freedom is preferable to confinement: he has most reason to prefer freedom, even though he has most reason to prefer to prefer confinement. Freedom is, in the sense I have in mind, what he ought to prefer, even though he ought to want to prefer confinement. Likewise, in the same sense of 'ought', the suspicious wife ought, perhaps, to believe that her husband is unfaithful, even if she ought to *want* to believe that he's faithful.² The correct belief, if all this is right, seems to be the one she ought, in this sense, to have. Or so I'll try saying, and see where it leads.

Beliefs are costly, and it would be silly to form sharp degrees of belief concerning everything, even if one could with enough cost and effort get them all correct. The sense of 'ought' I have in mind, though, ignores the costs of thinking. In this sense, it delivers ideal standards. Consider a normative theory of subjective probabilities or degrees of credence. One ought not to place high credence both in a proposition and its negation. The suspicious wife ought not, for instance, both to be pretty sure he's cheating and pretty sure he isn't. That would be incoherent. On many matters, to be

sure, ironing out incoherencies would be more trouble than it's worth. Costs counted, one ought to save mental effort and leave incoherent credences of little import undisturbed. The normative standards of subjective probability theory, though, ignore these costs, at least in the first instance. In the sense of 'ought' I mean to elucidate, one ought never to place high credence both in a claim and in its negation. The kind of *ought* I have in mind ignores the costs of thinking.

On the picture I am proposing, a single, primitive ought of this kind applies to a wide variety of states of mind: to action and intention, to wants and preferences, to feelings like admiration, and to beliefs and degrees of credence. The costs of thinking don't bear on the degrees of credence one ought (in this primitive sense) to have. They do bear, to be sure, on which ways of being more or less confident of claims are preferable. Still, what's preferable is, by the meaning of the term, what one ought to prefer costs aside. The costs of preferring freedom to confinement don't bear on whether freedom is preferable to confinement. Likewise, what's admirable is what one ought to admire—not what it's preferable to admire. Epistemic oughts, I'll try saying, are just this ought applied to belief and degrees of credence. What's most credible is a matter of which alternative one ought, in this sense, to place most credence in.

This brings us to an important contrast in kinds of ought. We can ask what one ought to do in light of all the facts. Alternatively, we can ask what one ought to do in light of available information. Late for an important meeting, I approach a blind intersection. In fact, nothing is coming on the crossroad, and so in light of all the facts, I ought to drive on through without slowing down. I have no way of knowing this, however, until I slow down and look, and so in light of my information, I ought to slow down and look, and proceed only if I see that nothing is coming. Standardly in moral theory, we distinguish what a person ought to do in the *objective* sense and what she ought to do in the *subjective* sense. Or the tradition speaks of right and wrong acts, and distinguishes objective and subjective senses of these terms.³ An objectively right act we might more aptly call, following Bertrand Russell, a morally "fortunate" act. The same distinction applies to the more primitive, non-moral ought that I'm exploring.⁴

The distinction applies to oughts of belief. You flip a coin and hide the result from both of us. If in fact the coin landed heads, then in the objective sense, I ought to believe that it landed heads. Believing the coin landed heads would be, we might say, epistemically fortunate. In the subjective sense, though, I ought neither to believe that it landed heads nor believe that it landed tails. I ought to give equal credence to its having landed heads and to its having landed tails. The coin in fact landed heads, imagine, and so the correct belief for me to have is that the coin landed heads. Such a full belief would be silly, but it would be correct. The correct belief, then, we can try inferring, is the belief one ought to have, where this cost-ignoring *ought* is in

the objective sense. The correct belief that the coin landed heads, though, is not the one I ought in the subjective sense to have; I ought subjectively to give equal credence to its landing heads and its landing tails.

In the objective sense, then, one ought to believe that the coin landed heads just in case, in fact, the coin landed heads. From an *is*, that the coin landed heads, we get an objective *ought*, that I ought objectively to believe that the coin landed heads. The coin's landing heads isn't merely a supporting ground for believing that it landed heads, as rain is a supporting ground for taking an umbrella. The coin landed heads, and from this, it seems, follows analytically a normative conclusion: that objectively, I ought to believe that snow is white.

I would like to explain why this is analytic—and explain it without denying a sharp is-ought distinction. My proposal, in outline, will be this: The subjective ought is the primitive one, and the objective ought can be defined in its terms. Once we devise a way to do this, it will fall out why, using this definition, we can derive an objective ought from the coin's landing heads. That one objectively ought to believe that the coin landed heads is, it will turn out, a normative claim only degenerately. When we expand the claim in terms of subjective oughts, the oughts in effect cancel out. That is the story I'll tell.

1. Alternative Ways Out

One candidate way out of the puzzle I have raised is proposed by Paul Boghossian and Nishi Shah.⁵ Perhaps the concept of belief itself is normative. "Belief", we can try saying, is that propositional attitude which is correct, toward *S*, just in case *S* is true. It then becomes trivially analytic that one ought to believe *S* just in case *S* is true. That's just what the word 'believe' means.

Clearly, if that's what 'believe' means, then this proposal works. I find, though, that I can't accept that the proposal gives what we mean by the word 'believe'. Surely what happens can't be this: I find myself in a state with regard to the content that snow is white, and ask myself, "Can this be belief?" I then ask myself whether it's a state that's correct just in case snow is white. Finding that I answer yes, I conclude that it indeed is belief. Such a procedure works well for bachelorhood: Knowing that John is a grown man, I inquire whether he is married, and on learning that he isn't, I conclude that he is a bachelor. Such a procedure works well for love: observing a glow in my chest whenever I see her, I ask "Can this be love?" As I check of the fit of my symptoms with those of song and story, I conclude that it is.

I don't have a definitive argument against this kind of proposal. But, as Gareth Evans argued, I find out whether I believe that snow is white by

asking myself whether snow is white and answering yes.⁶ Why would this be a test for whether I'm in a state whose correctness condition is that snow be white? Perhaps an answer can be found, but it is worth finding an alternative explanation of why truth is the correctness condition for belief.

Alternatively, we might try construing the 'ought's in the statements that worried us as implicitly hypothetical: "If truth is to be the only object, then this is what one ought to accept."⁷ And in a way, indeed, that will be my explanation of the paradox. There's no special *is/ought* problem when the *ought* is hypothetical; hypothetical oughts can follow analytically from the facts. Think of *oughts* as equivalent to imperatives of a special kind; then hypothetical *oughts* are equivalent to hypothetical imperatives. And the validity of a hypothetical imperative can be fully a matter of fact. For example, from the *is* statement

If you were to climb out the window, you would escape the fire, and otherwise you wouldn't,

follows the hypothetical imperative,

¡If you want to escape the fire, climb out the window!. (1)

Hare has taught us to see a hypothetical imperative like (1) as a conditional with an imperatives both in the antecedent and in the consequent:

If ¡Escape the fire! then ¡Climb out the window!.

(I indicate imperatives with a fusion of German and Spanish language punctuation.) This kind of imperative does follow from an *is*. Anyone who accepts the facts must accept it, since sufficient norms to apply, given the facts, are introduced hypothetically in the antecedent. The logic of the concepts involved requires accepting the inference from the *is* premise to the hypothetical imperative conclusion no matter what substantive normative views are to be accepted.

If, then, the *oughts* of correct belief are hypothetical in this way, that might explain why some of them are equivalent to non-normative statements. Take this line, and we then don't need to give up an *is/ought* gap in general. We can keep it and still admit that the form of a complex, hypothetical *ought* statement might make it follow from an *is*.

We now do get that the hypothetical imperative treatment of norms for belief works. If snow is white and you want a true belief on the matter, then believe that snow is white. That's truistic, though it puts the truism misleadingly. More precisely, if snow is white, then if you ought to believe truly on the matter, then you ought to believe that snow is white.

All this seems undeniable—but surely it’s not enough. True belief is correct, not just correct given the correctness of true belief. We need to explain how this can be, when correctness is a normative notion. It is not satisfactory just to declare these *oughts* hypothetical, in an *ad hoc* way, and declare matters resolved. What is it about the ‘ought’s in our puzzle, we need to ask, that means we can treat them hypothetically?

2. Objective ‘Ought’s: A Strategy

Whether an ought is objective or subjective is not a matter of how objective the ultimate normative basis for the judgment is. Ultimate norms for acting on limited information might be objectively based in whatever sense principles for acting in light of all the facts turn out to be so. The distinction is a matter of how much by way of the non-normative *facts* one needs to know to ground a non-ultimate *ought*. The basic normative precepts that ground a subjective *ought* are subjectively applicable—applicable in light of information the agent has. The basic precepts that govern objective *oughts* are not. They are objectively applicable: their application may require access to facts the agent has no way of ascertaining.

Consider a hedonistic utilitarian of Sidgwick’s kind.⁸ He thinks it rational to promote pleasure in the universe, irrespective of whose it is. If, now, he tells us to maximize net pleasure in the universe, the precept he’s giving us is applicable objectively. Normally, we won’t have the information we need to guide ourselves by it. Suppose, though, he says to maximize one’s *subjectively expected value* of net pleasure in the universe. This basic precept is subjectively applicable. Applying it may, to be sure, require superhuman powers of calculation, but sheer lack of information won’t by itself keep one from complying.

Now in a way, objective *oughts* seem fishy. It would be *nice*, to be sure, to know what objectively you ought to do. If indeed you could check with an omniscient advisor, that’s what you would ask: If, say, you wonder whether to take an umbrella, who better to ask than someone with detailed foresight of whether it will be raining at moments you’ll want to be outside? Such advisors, though, are in short supply. Your real questions, then, are what to do on the basis of information you have.⁹ (This includes whether to take steps to extend your information. Information has its costs, at least in time and effort, and so seeking more information is sometimes wise and sometimes not.) The *ought*-precepts you need, we can argue, will be for ‘ought’ not in the fishy objective sense, but in the *subjective* sense: in light of what you know.

The *oughts* that raised problems for us are not only objective, but in a broad sense *epistemic*: What they govern directly is *belief* or *acceptance*. These *oughts* govern not the will or the wish to believe, but belief itself. Now

with epistemic *oughts*, an objective sense will be especially strange. These *oughts* tell us what to believe, but tell us what to believe in light of things we have no way of knowing. An *ought* is supposed to be normative; it is meant to guide us. With these *oughts*, though, we may not know what we'd need to know in order to recognize which ones to accept. This, perhaps, will explain why some of these objective, epistemic *oughts* have the property at the root of our worries: analytical equivalence to non-normative statements of fact.

Here, then, is a strategy of investigation: Take subjective *oughts* as basic, and explain objective *oughts* in terms of them. Subjective *oughts*, I'll claim, can't be dismissed as somehow hypothetical or covertly factual. These are the *oughts* that exert normative governance—the *oughts* we accept and whose acceptance is directly motivating.¹⁰ We must act, after all, in light of information we have. Subjective epistemic *oughts* in particular tell us how evidence is to be assessed, and disagreements over evidence are genuinely normative: they are disagreements as to what to take as grounds for believing what. Here, then, is the strategy: Take subjective *oughts* as basic, and look for a general way to explain objective *oughts* in terms of subjective *oughts*. Then apply the general explanation to the special case of epistemic *oughts*. What happens, we can then ask, to the analytic entailments that had threatened a sharp fact/norm distinction? The *oughts* that can't follow from *iss*, I'll guess at the outset, are subjective. That leaves it open whether an *ought* can follow from *is* when the *ought* is objective.

In the epistemic case, I've been suggesting, objective *oughts* are hypothetical in a way that makes them not genuinely normative. If so, I'm now saying, this conclusion should follow from more general principles—and I'll be claiming that it does. In the end I'll conclude that objective *oughts* are indeed hypothetical. If, then, an objective *ought* follows from a non-normative *is*, that raises no problem for a sharp fact/norm distinction. It's fairly uncontroversial, after all, that an imperative can follow from a matter of fact if the imperative is only hypothetical. Our program, then, will be to see in what way objective *oughts* are hypothetical, and how this makes certain objective *oughts* equivalent to non-normative facts.

3. Characterizing Objective 'Ought's

I'll first review objective and subjective *oughts* in general, and ask more systematically whether either can be characterized in terms of the other. I'll endorse in the end what I've already been asserting, that subjective *oughts* are basic. Objective *oughts* can be defined in terms of them, I'll conclude, and subjective *oughts* can't be defined in terms of objective *oughts*.

First, this last point: We can't define a subjective *ought* in terms of the objective *ought*. We might try saying that an act is right in the subjective sense just in case it would have been right in the objective sense if the facts

had been as the agent thought them to be.¹¹ Often, though, an agent will know himself ignorant, and in that case there is no way he thinks the facts to be. Alternatively, we might try saying that an act is subjectively right just in case it would have been objectively right if the facts had been as the agent thought them *most likely* to be. If an act risks disaster, though, this may make it prospectively worth avoiding even though the agent thinks it likely beneficial. At a blind intersection, for instance, speeding on through is subjectively wrong, even if most likely no cross-traffic is coming.

It should be easy to see, indeed, that no definition of the subjective sense in terms of the objective sense will work. Suppose you are offered a bet on the flip of a coin: win a thousand dollars if the coin lands heads, and lose \$800 if it lands tails. You have no way of knowing how the coin will land. To judge what you ought to do objectively, we need to know how the coin will in fact land, but we don't need to know much about the value of money: We just need to know that having more is better than having less. To decide what you ought to do subjectively, in contrast, we need to know more about value: how the gain of a thousand dollars compares to the loss of \$800. In terms of classical decision theory, the utility needed for objective *oughts* is ordinal, whereas that for subjective *oughts* is cardinal: the scale must allow for comparisons of preferability in degrees. The argument, then, that we can't characterize subjective *oughts* in terms of objective *oughts* boils down to this: from objective *oughts* we can glean only an ordinal utility scale for the sure alternatives. What one ought to do subjectively depends not only on this, but on the cardinal utilities involved. The onlooker, then, who knows objective *oughts* doesn't thereby have the information he would need to settle what the subject ought to do subjectively.

Likewise with belief: The facts settle straightforwardly what one ought to believe objectively. One ought to believe all truths and disbelieve all falsehoods—with certainty. Much more is needed to settle what degrees of credence one ought subjectively to have. If the subjective ought could be defined in terms of the objective ought, then objective oughts would settle, analytically, what one's degrees of credence ought subjectively to be. But that's a question of how to assess evidence. How could someone who knows all the facts determine, from sheer analytic definitions, how, subjectively, I ought to assess my evidence?

The question, then, is whether we can go the other way, and characterize objective *oughts* in terms of subjective *oughts*. In this direction things work better. What ought one objectively to do? As a rough gloss, we might try this: *It is what it would be the case that one ought subjectively to do if one had full information*—if one learned everything that is the case.

The syntax here is awkward: The verb 'ought' lacks the form I needed, and so I had to make do with longwinded talk of "what it would be the case that one ought" to do. The problem of missing forms of 'ought' will recur, and so I need to find a way to evade it. I'll resort, then, to English

supplemented by an ugly stipulation. I'll use 'to ought' as an infinitive of a regular verb, so that instead of "It would be the case that I ought to" I'll say "I would ought to". Instead of "If it were the case that I ought to", I'll say "If I oughted to".

Rephrase the proposal, then: I'll use a subscript for the objective sense, so that "ought_{ob}" means ought in the objective sense, and leave the subjective sense unmarked. An unmarked 'ought', then, will be in the subjective sense unless otherwise indicated. Then, we can say, what one *ought_{ob}* to do is what one would ought to do if one had full information. It's what it would be the case that one ought to do, in the subjective sense, with full information.

This is the proposal in rough, but glaringly it requires work. How shall we render "full information"? To have full information is to know everything that's the case—but what does this mean? That I believe the full truth? So we might think, but it won't do: Information or knowledge takes not only true belief, but true belief that is justified. Worse still, true belief by itself won't affect what one *ought* to do—or won't affect it in the right way. Suppose I have an unjustified belief *N* as to the exact number of eggs in Ann Arbor, and fortuitously, *N* happens to be true. I believe statement *N*, but I ought not to believe it. Since, then, I ought not to believe *N*, I ought not to act on my belief in *N*. If offered a bet on *N* at even odds, I'd ought to decline the bet, since I have no reason to think *N* true and every reason to think it highly unlikely.

What I ought to do in light of full information, then, is not what I'd ought to do if I believed the whole truth with no justification. It is what I'd ought to do if I believed the whole truth *with* justification. Indeed it wouldn't matter, for this purpose, if I didn't believe it, so long as I *oughted* to believe it—so long as this were what I *ought* to believe. What I ought to do, in light of my epistemic situation, is a matter of what I ought to believe—it's a matter of what I ought to do in light of the beliefs that I ought to have.

To say that I ought_{ob} to do *A*, then, is to say this:

If I oughted to accept all the facts, I'd ought to do *A* (2)

What is it, then, to accept all the facts? One can't do it simply by saying to oneself "I hereby accept all the facts." One must accept a full and true factual description of one's circumstances.

This leads to an crucial subtlety in interpreting a conditional like (2). One's circumstance will include much ignorance—but if one knew everything, one wouldn't be ignorant. One then wouldn't need to acquire books, newspapers, almanacs, and the like or do internet searches. Often, though, one ought_{ob} to acquire sources of knowledge; they'll guide one later on. In interpreting (2), then, we'll have to be careful about which facts I'm to ought

to accept. When I say “If I oughted to accept all the facts”, I mean not to the facts that *would* obtain under that wild supposition, but the facts that *do* obtain—my ignorance included.

I’m a spy in a good cause, imagine, behind enemy lines. I need to report the enemy’s assembled strength. Ought I, objectively, to risk my life to observe it? If I oughted to accept all the facts, that would mean I already had a reliable source of information, and didn’t have to look. I might, say, have been whispered the figures on assembled strength by a soothsayer of proven track record. Risky observations then would be superfluous and foolhardy. The phrase ‘all the facts’ in (2), though, includes not only enemy strength, but my ignorance of it. Call myself if I had access to all the facts I^+ . I^+ ought to accept all the facts not of his own fanciful circumstance, but of my actual circumstance. My actual circumstance includes the enemy’s strength, but also my ignorance of it—and the facts in virtue of which I need to know it.

This makes I^+ a strange character: he ought both to believe the enemy’s strength to be what it in fact is, and to believe himself ignorant of the enemy’s strength. It’s hard to see what circumstance, however fanciful, could put me in such a frame of mind. A better way to think of the matter might be this: I^+ decides matters not for himself but for me—for my circumstance. The question is what I^+ ought to decide for my circumstance, with all my ignorance. To say that objectively, I ought to do *A*, then, is to say something like this: Suppose a duplicate of mine, I^+ , were given all the facts of my situation, and were deciding, for my circumstance, what to do. Then I^+ ought—subjectively—to decide, for my circumstance, to do *A*. Knowing, say, that the enemy has assembled five depleted regiments, that I’m ignorant of this fact, and that the place where I’d sneak up and look is unguarded, I^+ ought to decide, for my circumstance, to sneak up and look.

That I *ought*_{ob} to do *A*, we’re now saying, means this:

If a duplicate of mine I^+ were transformed so that he oughted to accept all the facts of my circumstance, and I^+ were to decide, for my circumstance, what to do, then I^+ would ought to decide, for my circumstance, to do *A*. (3)

What kind of supposition is I^+ ’s omniscience? We’re supposing that I^+ were to ought to accept everything that holds for my actual circumstance.

Clearly the supposition is wildly counterfactual. How might it be that I^+ oughted to accept all these things that I myself have no way of knowing? His factual circumstances would have to be quite different from mine: Perhaps he’d have been whispered everything by an all-seeing soothsayer, and have unlimited powers of memory. Perhaps he’d have to be all-seeing himself.

Is a conditional like (3) counternormative as well as counterfactual? I’ll leave this question open: I’ll take it that the wild supposition of (3) is sure to

be counterfactual, and allow that it might be counternormative as well—but may not be. Perhaps, given the right epistemic norms, we can make it the case that I^+ ought to accept all the facts by giving him access to an all-seeing soothsayer, along with sufficient demonstrations of the soothsayer's prowess. Perhaps we can make I^+ himself all-seeing. What, though, if the right epistemic norms are more skeptical: they forbid full credence in the whisperings of soothsayers, no matter how impressive their demonstrations of sooth-saying prowess. They forbid full trust of what one appears to see. In that case, the wild supposition of (3) must change not only facts about I^+ , but the norms that validly govern him.

Ideal judge I^+ , remember, is a person who ought to accept all that's the case for my actual situation, not for his. He is me as I would be were it the case that I ought to accept all that, as things are for me in actuality, is the case. He expects forthwith to be me in my situation, with all my limitations, and his thoughts concern the situation he expects forthwith to be in. Since I am ignorant, he ought to accept that he is to be ignorant. He ought to accept the facts in virtue of which I am ignorant, such as not being able to see down the crossroad. He ought to accept the oughts that apply to me in my actual situation, not those that apply to him. Thus, for anything S that is the case for me now—such as that nothing is coming on the crossroad but that I can't see that nothing is coming— I^+ ought to accept that S . This includes oughts that apply to me. Suppose I ought not to believe what I can't see to be the case. Then I^+ ought to accept that he ought not to believe what he can't see to be so. The situation of I^+ , though, as opposed to the one he thinks himself to be in and is accepting norms for, may be counternormative. Suppose that it is; suppose he ought to accept this: that nothing is coming although that's not something he can see. (Perhaps he feels a strong intuition, these intuitive feelings are accurate, and counternormatively, he ought to accept the deliverances of his strong intuitive feelings. Or perhaps, he has a crystal ball, what he sees in it accurately mirrors what is the case from my perspective, and counternormatively, he ought to accept whatever the crystal ball depicts.) He ought to accept that nothing is coming, even though he ought also to accept, for the state he expects forthwith to be in, that he ought to suspend judgment on whether anything is coming. He ought to believe what he can't see to be the case, though he also ought to accept, for the state he expects forthwith to be in, that he ought not to believe what he can't see.

4. Correct Belief

We now have an informal account of how objective *oughts* work in general—how they can be defined in terms of subjective *oughts*. What, then, of the objective *oughts* that raised our problem in the first place? Where S is

the content of a statement, our puzzle was, the objective norm i_{ob} . Accept that $S!$ seems analytically equivalent to S . This seemed, then, to be a case of getting a normative *ought* from a non-normative *is* analytically. My hypothesis was that once we characterized objective *oughts* in terms of subjective ones, we'd see why the equivalence was analytic—and in addition, that we would see that there is still a gap between an *is* and *ought*, when it is the kind of *ought* that is basic, the subjective *ought*. Does this follow from what we've been saying?

Our problem, recall, concerns not statements as linguistic utterances, but the contents of statements: the thought, for instance, that snow is white. (That the English sentence “Snow is white” is correct is a fact not only about snow and its color, but about the English language.) Let that S be the content of a factual statement: that S , for instance, might be that snow is white.

Turn now to the characterization of objective *oughts* in (3). The wild supposition of (3) was that I^+ oughted to accept, for my circumstance, all the facts of my circumstance. I^+ , then, would ought to accept that S if and only if S obtains—in my circumstance. To say that I *ought_{ob}* to accept that S is to say that I^+ ought, for my circumstance, to accept that S . The analytic equivalence we were seeking therefore holds. We have:

I *ought_{ob}* to accept that S iff I^+ ought to accept that S for my circumstance. (4)

I^+ ought to accept that S for my circumstance iff S obtains in my circumstance; (5)

(4) is our proposed definition of the objective *ought*, and (5) is the characterization of I^+ . I^+ , after all, is just a duplicate of me, transformed—factually and perhaps normatively—so that (5) holds for all S . From (4) and (5) it follows that I *ought_{ob}* to accept S if and only if S obtains in my circumstance. Truth is the condition for correct belief.

Notes

1. Rosen, in “Meaning, Normativity” (1999), is cited and quoted in Shah, “How Truth” (2003), 458. Calling a performance of Jingle Bells “correct” is not “to say that it was played as it ought to be played (Shah, 458). Boghossian, in “Normativity of Content” (2003), 37, also denies that correctness facts are ought facts.
2. This distinction is widely noted. For this diagnosis of the distinction, see my *Wise Choices* (1990), 37.
3. See, for instance, Brandt, “Toward a Credible”, 110–115. I discuss the two senses in *Utilitarianism and Coordination* (1990), 29–36.

4. I propose in *Wise Choices* (1990) that terms like ‘right’, ‘wrong’, and ‘ought’ in their moral senses can be defined in terms of this primitive ‘ought’. Both *Wise Choices* and *Thinking How to Live* develop a theory of what this primitive ‘ought’ means.
5. Boghossian, “Normativity of Content” (2003), 40; Shah, “How Truth Governs Belief” (2003), 465–474. Boghossian offers this diagnosis as an alternative to the frequent claim that the concepts of meaning and of mental content are normative concepts. I discuss the “normativity of content”, inconclusively, in “Meaning and Normativity” (1994), “Thought, Norms” (1996), and “Thoughts and Norms” (2003).
6. Evans, *Varieties* (1982), 225.
7. Horwich, *Truth* (1990, 1998); “Norms of Truth” (2001).
8. Sidgwick, *Methods of Ethics* (1907). Sidgwick himself uses ‘ought’ only in the objective sense.
9. I discuss this in *Utilitarianism* (1990), 37–42.
10. I speak of “normative governance” in *Wise Choices* (1990), 72–77.
11. Brandt, in “Toward a Credible” (1963), considers such a definition, 113–114. I discuss such definitions in *Utilitarianism and Coordination* (1990), 31–34.

References

- Boghossian, Paul (2003). “The Normativity of Content”. *Philosophical Issues* 13:1 (Oct.), 31–47.
- Brandom, Robert (1994). *Making It Explicit* (Cambridge, MA: Harvard University Press).
- Brandt, Richard (1963). “Toward a Credible Form of Utilitarianism”. Hector-Neri Castañeda and George Nakhnikian, *Morality and the Language of Conduct* (Detroit: Wayne State University Press), 107–143.
- Evans, Gareth (1982). *The Varieties of Reference* (Oxford: Oxford University Press), 225.
- Gibbard, Allan (1990). *Utilitarianism and Coordination* (New York: Garland Publishing).
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Oxford: Oxford University Press).
- Gibbard, Allan (1994). “Meaning and Normativity”. *Philosophical Issues* 5, Enrique Villanueva, ed., *Truth and Rationality* (Atascadero, CA: Ridgeview Publishing Co., 1994).
- Gibbard, Allan (1996). “Thought, Norms, and Discursive Practice: Commentary on Robert Brandom, *Making it Explicit*”. *Philosophy and Phenomenological Research* 56, 699–717.
- Gibbard, Allan (2003). *Thinking How to Live* (Cambridge, MA: Harvard University Press).
- Gibbard, Allan (2003). “Thoughts and Norms”. *Philosophical Issues: A Supplement to Nous* 13:1, 83–98.
- Horwich, Paul (1990). *Truth* (Oxford: Oxford University Press). Second edition, 1998.
- Horwich, Paul (2001). “Norms of Truth and Meaning”. R. Schantz, ed., *What Is Truth?* (Berlin: de Gruyter).
- Rosen, Gideon (1999 unpublished mss.). “Meaning, Normativity, and All That”.
- Sidgwick, Henry (1907), *The Methods of Ethics*, Seventh Edition (London: Macmillan).
- Shah, Nishi (2003). “How Truth Governs Belief”. *Philosophical Review* 112, 447–482.