

What You Ought to Believe
Between Truth and Justification
John MacFarlane and Niko Kolodny (UC Berkeley)

1. INTRODUCTION

Can we say anything general about what one ought to believe? We are driven between unpalatably subjective and unpalatably objective extremes.

Justification Norm If S is (or ought to be) considering p , S ought to believe that p iff S has justification for believing p .

Truth Norm If S is (or ought to be) considering p , S ought to believe that p iff p is true.

The Justification Norm is too subjective; it can't explain uses of "ought" in **Advice**:

Ignorant Toss. A coin is tossed. Informed peeks and sees that it landed heads, but Ignorant does not see the coin.

IGNORANT: I ought not believe that it landed heads, and I ought not believe that it landed tails.

INFORMED: No, you ought to believe that it landed heads.

The Truth Norm is too objective; it demands omniscience, and is of limited use in guiding belief formation. On this view, it would never be true that one ought to suspend belief (believe neither p nor not- p).

One might try explaining Deliberation with the Justification Norm and Advice with the Truth Norm:

Ambiguity Theory Some uses of "ought" are "objective" (governed by the Truth Norm) and others are "subjective" (governed by the Justification Norm).

But this has Ignorant and Informed talking past each other. On this view, it's hard to see how Informed's advice is relevant to the question Ignorant was asking. It's also hard to see why Ignorant should *retract* her assertion after learning from Informed that the coin has landed heads.

Moreover, Ambiguity Theory won't handle all cases:

Deceived Toss. Deceived believes, with justification, that the coin is a trick coin, with heads on both sides, while Undeceived knows that the coin is fair. Neither sees how the coin landed.

DECEIVED: I ought to believe that it landed heads.

UNDECEIVED: No, you ought not believe that it landed heads, and you ought not believe that it landed tails.

Undeceived can't be using "ought" in either the objective or the subjective sense.

It seems that what we need is something **between truth and justification**. What about *knowledge*, which is connected to reasons and evidence (like justification) but also to truth?

Knowledge Norm If S is (or ought to be) considering p , S ought to believe that p iff such a belief would constitute knowledge that p .

This handled Deceived Toss well: Deceived thinks she ought to believe that p because she thinks that such a belief would constitute knowledge. Undeceived thinks Deceived ought not believe that p , because he knows that this belief is not knowledge (since it is based on misleading evidence).

But the Knowledge Norm doesn't handle Ignorant Toss any better than the Justification Norm.

2. AN ALTERNATIVE

Likelihood Norm If s is (or ought to be) considering p , S ought to believe that p iff p is (sufficiently*) likely. [Aside: This need not be a threshold view. If you don't think we should believe that lottery tickets will lose, you can understand "sufficiently likely" as "certain," or you can make "sufficiently" depend on practical circumstances.]

The Likelihood Norm can make sense of Ignorant's assertion, Informed's reply, Deceived's assertion, and Undeceived's reply. What about Advice?

Start by considering standard contextualist lines on the meaning of "likely":

Nonrigid Contextualism '*Likely ϕ* ' is true at c, w iff it is likely, in light of the information possessed in w by the relevant group at c , that ϕ is true at c, w .

Rigid Contextualism '*Likely ϕ* ' is true at c, w iff it is likely, in light of the information possessed in w_c by the relevant group at c , that ϕ is true at c, w .

Neither adequately explains Advice. If Ignorant and Informed are not talking past each other, their "oughts"s must be evaluated with respect to the same contextually relevant group. If Informed's remark is to be apt, that group must include Informed. But then Ignorant's assertion is unjustified, since Ignorant knows that Informed peeked at the coin.

Group contextualism also loses the **transparency** of doxastic deliberation. The question whether I ought to believe that p should bottom out into a question whether p . But on this proposal, it bottoms out into a question about what my fellow group members know.

But fortunately for the Likelihood Norm, there are independent reasons to reject Contextualism about "likely," and adopt a different, "relativist" semantics. When "likely" is understood correctly, I will claim, the Likelihood Norm does explain Advice.

The argument has two steps. First, reasons for thinking that '*Likely ϕ* ' does not express a proposition about any group's information. This establishes that "likely" is not indexical. Second, reasons for thinking that it is the assessor's information, not the speaker's, that matters for the accuracy of a likelihood claim. This establishes that "likely" is assessment-sensitive.

3. AGAINST INDEXICALISM

- ★ If we counterfactually vary only one's information, but not the facts, then what is likely does not vary.
 - (1) If it were still the case that the butler had opportunity, but I had no evidence of this, it would still be likely that he did it.
 - (2) # If it were still the case that the butler had opportunity, but I had no evidence of this, it would not still be likely that he did it.

This speaks against Nonrigid Contextualism.

- ★ If we vary the facts, without varying one's information, then what is likely varies:
 - (3) If the butler had not had opportunity, but my evidence still indicated that he had, then it would not still be likely that he did it (though I would wrongly think that it was).
 - (4) # If the butler had not had opportunity, but my evidence still indicated that he had, then it would still be likely that he did it.

This speaks against both Rigid and Nonrigid Contextualism.

- ★ Likelihood claims behave differently from claims about what one's evidence makes likely, when embedded in the antecedents of indicative conditionals (Yalcin):
 - (5) If not- p and my evidence makes p likely, then ...
 - (6) If not- p and my actual evidence makes p likely, then ...
 - (7) # If not- p and it is likely that p , then ...

In the last case there seems to be no coherent hypothesis to entertain. This tells against both Rigid and Nonrigid Contextualism.

Nonindexical semantics for "likely" '*Likely ϕ* ' is true at c, w, i iff the information i makes it likely that ϕ is true at c, w, i .

On this view, a likelihood claim is not *about* any particular body of information, any more than a contingent claim is *about* any particular world. Just as a contingent claim has truth values relative to worlds, so a likelihood claim has truth values relative to (worlds and) information states.

4. IN FAVOR OF RELATIVISM

The nonindexical semantics is consistent with a **nonindexical contextualist** view, which takes the accuracy of a likelihood claim to depend on its truth relative to the information relevant at the context of use, and also with a **relativist** view, which takes the accuracy of a likelihood claim, as assessed from some context, to depend on the information relevant at that context, the context of assessment. So why prefer the relativist version?

The relativist version better captures our sense that Informed is *disagreeing* with Ignorant. For only the relativist version predicts that Ignorant should feel pressure to *retract* her original claim after gaining new information from Informed. This is what we need to explain Advice.

Construed this way, the Likelihood Norm can explain what was appealing about both the Truth Norm and the Justification Norm.

- ★ The Justification Norm is appealing, in part, because it feels incoherent to say:

(8) # I am not justified in believing that p , but I ought to believe that p .

The Justification Norm would explain why this is incoherent. But so does the Likelihood Norm, since “it is likely” and “I am justified in believing it” coincide from a first-person present point of view.

From a third person point of view, however, it makes perfect sense to say:

(9) He is not justified in believing that p , but it’s likely that p .

(10) He is not justified in believing that p , but it’s what he ought to believe.

That’s what an advisor might think before giving advice.

- ★ The Truth Norm is appealing, in part, because it seems incoherent to say:

(11) # He ought to believe that p (that’s what I’d advise him to believe, if I could), but p is false.

The Truth Norm would explain why this is incoherent. But so does the Likelihood Norm, since the following is also incoherent:

(12) # It is likely that p , but p is false.

Another fact that seems to favor the Truth Norm is the plausibility of conditionals with the following form:

(13) If p , then S ought to believe that p .

(14) If not- p , then it is not the case that S ought to believe that p .

The Truth Norm can explain this, taking the conditionals to be material conditionals. But the Likelihood Norm can also explain it, if we take the conditionals to be indicative conditionals, and analyze them as in “Ifs and Oughts,” as information state shifters.

5. GIBBARD’S PUZZLE

This yields a nice resolution of a puzzle posed by Gibbard in “Truth and Correct Belief.” Gibbard suggests that one can infer ‘You ought to believe p ’ from p , and worries that this looks like an example of an analytic entailment of a normative claim by a purely descriptive one. His rather complex solution to the puzzle is to offer a definition of the objective “ought” in terms of the subjective “ought” and counterfactuals. Given this definition, the entailment in question is analytic,

but, Gibbard thinks, this is just an artefact of the way the objective “ought” is defined, and does not reveal a genuine counterexample to the Humean dictum that you can’t derive an *ought* from an *is*.

We think that our account provides a simpler analysis of the puzzle. What makes it seem as though there is an analytic entailment from p to ‘ S ought to believe that p ’ is, first, that the conditional, ‘If p , then S ought to believe that p ’ is logically true and, second, that when p is *known*, one can infer ‘ S ought to believe that p ’. But since modus ponens does not preserve truth when the consequent is information-sensitive, there actually is no analytic *entailment* from p to ‘ S ought to believe that p . One cannot derive this ought from this is.

6. CRITICAL VS. DELIBERATIVE

The Justification Norm is supported, in part, by reflection on criticism. Feldman: “To say, for example, that ancient school children should have believed that the Earth was (roughly) round, since it is roughly round, is to miss the point of epistemic obligation altogether. By believing that the Earth is flat such children were believing exactly what they should have believed given the situation they were in.” (“Epistemic Obligations,” 246) The intuition, we think, is that it would be wrong to criticize the children; they were reasoning and assessing information admirably.

1. A believer is not criticizable when he believes what he is justified in believing, and is criticizable when he does not believe what he is justified in believing.
2. What someone is not criticizable for doing, and is criticizable for not doing, is what he ought to do.
3. Therefore, what a believer ought to believe is what he is justified in believing.

We think the second premise is mistaken. It ties “ought” too closely to criticism. The question of what I ought to believe is at root a *deliberative* question, characteristically asked in the first person (though others can take it up on the agent’s behalf, as in advice). The question whether a believer is criticizable for what he believes, by contrast, is a *critical* or *evaluative* one, characteristically asked in the third person. Such questions seek to assess the believer, to determine whether he should be praised or blamed, approved of or disapproved of, followed or shunned. (I *can* ask a critical question in the first person—but I’m adopting the viewpoint of a future spectator.)

“Justified” is chiefly critical and evaluative. When the question is whether we are justified, others’ information is irrelevant. But “ought” is chiefly deliberative. And when we deliberate, we want access to our potential advisers’ better information. We want to do what is really best, not just to do something later observers will judge as reasonable. That is, we think, why “ought” is assessment-sensitive, and “justified” is not.

We concede that sometimes “ought” is used in a critical or evaluative way:

- (15) Don’t beat yourself up over it. You believed then just what you ought to have believed.

We think this can be handled without positing an ambiguity, because of the flexibility of “information relevant at the context of assessment.” When the assessor is interested only in critical appraisal, rather than advice, the information relevant to the assessor will be the believer’s information at the relevant time.