

Against the Truth Norm*

Abstract

A variety of widely-held epistemological theories are committed to one or more of the following theses: that you objectively epistemically ought to believe all and only truths, or have maximally accurate credences; that the epistemic value of a person's beliefs is determined by their distance from the truth; and that rationality roughly consists in minimizing this distance to the extent possible, given one's limited information. These commitments form the foundation of accuracy-based epistemology and veritist epistemology generally. This paper focuses on the first claim. Call this the 'Truth Norm.'

This paper argues against the Truth Norm. The norm is open to a variety of interpretations, and each interpretation, once made precise, yields unacceptable verdicts about what we ought to believe. The worries I raise concern higher-order beliefs. Many epistemologists presuppose that puzzles generated by higher-order beliefs can be quarantined off, or idealized away, without impugning broad epistemological theories. One lesson of this paper: higher-orderdom is *infectious*. It pervades more of our beliefs—beliefs that matter to us—than is generally acknowledged. This fact will create problems for the Truth Norm and for many truth- and accuracy-based theories of epistemic rationality.

Epistemic rationality must have something to do with truth. Many epistemologists hold that norms of epistemic rationality—for example, prohibitions on incoherent beliefs¹—receive their force from their connection to truth. They're means to the end of having accurate beliefs. Philosophers in accuracy-first epistemology (AFE) hold that all genuine epistemic norms are just maximally efficient strategies for achieving accuracy (e.g., Rosenkrantz 1981; Joyce 1998; Greaves and Wallace 2006; Joyce 2009; Leitgeb and Pettigrew 2010a,b; Easwaran 2016; Pettigrew 2016; Briggs and Pettigrew 2018). Many epistemological theories are committed to one or more of the following theses: that you ought to believe all truths; that the epistemic value of a person's beliefs is determined by their distance from the truth; and that rationality roughly consists in minimizing this distance to the extent possible.

I focus on the claim that you ought to believe all and only truths. Call this the 'TRUTH NORM'. I focus on this claim because the tools of deontic logic turn out to be convenient for explicating the problem. Proponents of many of these theses, e.g. accuracy-first epistemologists, don't generally talk in terms of the objective epistemic *ought* most of the time. But they still claim that epistemic

*Thanks to Dan Greco, Ryan Doody, Melissa Fusco, Kevin Dorst, Brian Talbot, Bob Beddor, Simon Goldstein, Zach Barnett, Abelard Podgorski, Al Hájek, Nic Southwood, Joshua Schechter, David Christensen, Jamie Dreier, Miriam Schoenfeld, two anonymous referees, and audiences at the National University of Singapore, the Australian National University, Brown University, the Formal Epistemology Workshop, and the MITing of the Minds for valuable feedback.

¹I'll use 'belief' as shorthand for any doxastic attitude, including credence, suspension of judgment, etc.

rationality involves aiming to hold beliefs that best approximate an epistemic ideal: omniscience. Once we try to clarify what ideal we're meant to approach—what truths our beliefs should be aiming to approximate—we encounter the problems I raise.

This paper argues that the TRUTH NORM isn't a genuine epistemic norm. Why? Not merely because we ordinary believers can't satisfy it (in some sense of 'can't' that's sensitive to our cognitive limitations), or because it isn't action-guiding. Rather, the reason is that it's open to a variety of interpretations, and each interpretation, once made precise, yields unacceptable verdicts about what we ought to believe.

The worries I raise concern higher-order beliefs: beliefs wholly or partly about one's own beliefs. Many epistemologists presuppose that higher-order beliefs are just cute puzzle cases with no real probative force. They presuppose that we can quarantine off, or idealize away, worrisome higher-order beliefs. One lesson of this paper: there is no sharp line between first- and higher-order beliefs. Higher-orderdom is *infectious*. Many propositions we care about—e.g. whether a person is secretive, generous, responsible for deaths, rationally required to buy a lottery ticket, etc.—will turn out to be species of higher-order belief contents. And these will create problems for the TRUTH NORM, and for AFE in general.

The plan for the paper is as follows: in §1, I explain the notion of an objective epistemic 'ought'. I introduce a simple, schematic modal semantics for 'ought', and briefly discuss the ways in which it might imply 'can.' In §§2–4, I discuss three forms of interpretation of the TRUTH NORM that we can elucidate within modal semantics: *world-relative*, *absolute*, and *mixed*. I show that each form of interpretation generates unacceptable verdicts. In §5, I consider a weaker variant of the TRUTH NORM—the ACCURACY NORM—and show that it also makes unacceptable verdicts. I also show that its problems generate worries for a widely accepted form of AFE and other forms of epistemic consequentialism.

On this basis, I conclude: there is no TRUTH NORM on belief. There are no objective epistemic norms—at least, none of the sort that can play the roles in epistemology that an objective "ought" is expected to play. The central argument of this paper is openly inductive: I provide objections to several interpretations of the TRUTH NORM, starting with the simplest and most plausible and moving on to increasingly complicated and unintuitive variants. In §6, I discuss attempts to salvage some much weaker objective epistemic norms—norms that can't play the intended roles of the TRUTH NORM. I explain why even these should be rejected. At this stage, the search for any objective epistemic norm in the vicinity of the TRUTH NORM should start to feel like a degenerating research program. If such a norm exists, it won't be simple or intuitive. It will require giving up on widely used tools for understanding modality, e.g. Kripke- and Kratzer-style modal semantics. These considerations motivate skepticism about objective epistemic norms in the way that similar considerations have motivated skepticism about conceptual analyses of *knowledge*: the fact that it's so surprisingly hard to find an uncomplicated, intuitive statement of the thing is compelling (though obviously not decisive) evidence that it doesn't exist.

Finally, §7 explores an alternative conception of the evaluative landscape: the possibility of epistemic *value* without deontic implications, value that issues no objective 'ought'.

1 A Truth Norm

1.1 The objective epistemic ‘ought’

Philosophers often distinguish between subjective and objective ‘ought’s. The subjective ‘ought’ is sensitive to an agent’s limited information; the objective ‘ought’ is not. You *objectively* ought to take the very best option available to you, where what’s best is insensitive to your evidence. You *subjectively* ought to do what’s rational in light of your limited evidence.

A common thought in epistemology: the norms of epistemic rationality characterize the subjective *epistemic* ‘ought’.² They provide the best means of approximating the objective epistemic ‘ought’. The objective epistemic norm is widely thought to be the TRUTH NORM:³

TRUTH NORM: you ought to believe all and only truths.

For different belief models, the TRUTH NORM is better stated in other ways. If we (plausibly) distinguish between rejecting a proposition and believing its negation, then the TRUTH NORM should be: believe all truths and reject all falsehoods. If our belief model includes credences instead of, or in addition to, full beliefs, then the TRUTH NORM should require having credence 1 in all truths and 0 in all falsehoods. I primarily discuss full belief, but my arguments are generalizable to all three models.

The TRUTH NORM, on the intended interpretation, doesn’t say that you should *come to* believe all and only truths. It says that you should *already* believe all and only truths. The norm is best interpreted as time-coordinated: ‘you ought_t believe_t all and only truths’, rather than ‘you ought_t believe_{t+} all and only truths’. (Compare: while it’s true that you ought_t to pay_{t+} your parking tickets, it’s also true that you ought_t not have_t parking tickets to begin with.) This form of evaluative ‘ought’ isn’t action-guiding, since we can’t change the present; any change would take us into the future. But no one could reasonably expect an omniscience-requiring norm to guide actions.⁴

This paper ultimately argues that there is no objective epistemic ‘ought’. This needs explanation. There are vastly many possible forms of epistemic evaluation, many appealing to information outside the believer’s evidence. But that is not sufficient for labeling a form of evaluation ‘the objective epistemic “ought”’, on my usage.

In order for some norm to count as an ‘objective epistemic “ought”’, as I intend the phrase, it should be able to play certain roles in epistemology: for example,

1. It should characterize the epistemically *ideal* total and partial doxastic states.
2. It should provide the standard against which the epistemic value of other total doxastic states may be compared (as in AFE).
3. Divergence from the ideal should be able to affect (indirectly) whether a total doxastic state is rational. (If an objective epistemic ‘ought’ exists, it must be related to, and inform, the

²N.B. I use ‘epistemic “ought”’ for the *deontic* necessity modals used in epistemology. The expression is unfortunately ambiguous: ‘epistemic “ought”’ is sometimes used for (non-deontic) necessity given one’s knowledge.

³Some epistemologists take the objective epistemic ‘ought’ to concern knowledge rather than true belief. The objections that I’ll ultimately raise to the TRUTH NORM will apply equally to an analogous knowledge norm.

⁴I stipulatively define ‘omniscience’ in terms of true belief, rather than knowledge.

subjective epistemic ‘ought’. Note also that this role is part of why the objective epistemic ‘ought’ must be able to characterize *total* epistemic states.)

4. If a believer’s total doxastic state is as it objectively epistemically ought to be, then this doxastic state cannot be accuracy-dominated.⁵ (This is part of what makes a candidate objective ‘ought’ specifically *epistemic*, though it may not be sufficient.)
5. The objective epistemic ‘ought’ should not permit *hedges* against uncertainty, whereby one adopts a state that is known to be less than ideal in order to avoid the risk of a worse outcome. (Permitting hedging is sufficient for a form of evaluation to count as subjective.)

In principle one can slap the label ‘objective epistemic “ought”’ on ‘ought’s that cannot play most, or any, of these roles. But for the purposes of this paper, that will amount to changing the subject.

1.2 ‘Ought’ implies ‘can’

Some philosophers are skeptical of epistemic ‘ought’s in general. One motivation has been doxastic involuntarism (e.g., Alston 1988; for objections, see e.g., Chrisman 2008). Another motivation is skepticism about so-called “positive epistemic duties” (Wrenn 2007; Nelson 2010; for objections, see Ichikawa 2022, 2024; Simion 2023). Both styles of worry are fundamentally different from the worries for the TRUTH NORM this paper raises. Among other things, my objection targets specifically *objective* epistemic norms and is entirely friendly to *subjective* (evidence-sensitive) epistemic norms.

Another common objection to the TRUTH NORM comes from the charge that ‘ought’ implies ‘can’. Even setting aside doxastic involuntarism, there’s some sense in which we ordinary believers *can’t* believe all and only truths. Logical space includes uncountably many propositions: too many for cognitively limited believers like humans, with our little brains, to entertain. Moreover, some individual propositions might be impossible for us to cognize.

I’m not interested in this kind of objection. The objections to the TRUTH NORM that this paper raises, if correct, hold not only for ordinary believers but also for cognitively unlimited, ideal believers.

Even if we limit the set of propositions that the TRUTH NORM requires omniscience about, the problems I’ll raise remain: we needn’t assume that believers must have attitudes toward every proposition. Let \mathcal{F} be an agenda of first-order propositions that the agent can entertain. Assume that for all $p \in \mathcal{F}$, if a believer can entertain p , she can entertain Bp : the proposition that she herself believes that p . Let \mathcal{A} be the smallest boolean closed set of propositions that includes \mathcal{F} and also contains Bp for all $p \in \mathcal{F}$. The arguments below hold even if we assume that the agent can entertain only propositions in \mathcal{A} . So the believer might have only a small set of simple first- and second-order beliefs. Throughout I assume that the believer has attitudes toward some agenda of propositions $\mathcal{P} \supseteq \mathcal{A}$.

⁵Across epistemic possibilities, not merely logical possibilities.

1.3 Deontic logic for epistemology

What does it mean to say that you ought to believe all and only truths? Standard deontic logic (SDL) provides an off-the-shelf Kripke semantics for ‘ought’:⁶

- Propositions are sets of possible worlds.
- $[[\cdot]]$ is a valuation function assigning propositions to sentences. A sentence ϕ is true at w just in case $w \in [[\phi]]$.
- A *deontic accessibility relation* R is a serial relation over possible worlds.⁷
- v is a *deontic alternative* to w iff wRv for the relevant deontic accessibility relation R . Deontic alternatives to w are ideal worlds with respect to relevant norms in w .
- $[[\text{Ought } \phi]]$ is true at w iff $[[\phi]]$ is true at all deontic alternatives to w .

For ease of exposition, I’ll move freely between object and metalanguage, e.g., by re-expressing the previous sentence as ‘*It ought to be that ϕ at w iff at all deontic alternatives to w , ϕ .*’⁷

To interpret the TRUTH NORM, we need to locate an accessibility relation for its ‘ought’. In sections 2, 3, and 4, we’ll consider a series of candidates.

Note that, for all the talk of accessibility relations below, this central claims of this paper are *epistemological*, not logical. I will not argue that any possible interpretation of the Truth Norm will generate contradictions, or have Liar-like properties. Instead, I will argue that the natural and plausible interpretations are simply *false*. They make prescriptions that are not genuinely epistemically required, and more generally, they provide a form of evaluation that cannot fulfill the constitutive roles of an objective epistemic ‘ought’.

1.4 From deontic semantics to epistemic decision theory

AFE usually concerns decision theories, rather than modal semantics, for epistemology. These decision theories evaluate belief states in terms of gradational accuracy: for example, many hold that a belief state is irrational if it is strictly *accuracy dominated*, such that some other belief state is guaranteed to be closer to the truth, whatever the truth is. Since a major target of this paper is AFE, it’s useful to have a translation schema for epistemic decision theorists.

Each possible state of the world w determines at least one maximally accurate doxastic act. This serves various roles in AFE:

- A maximally accurate act at w is an ideal epistemic act at w .
- The accuracy or epistemic value of an epistemic act is measured by its proximity to a maximally accurate act.
- If your credences at w match a maximally accurate act at w , then your credences can’t be *strictly accuracy dominated*.

⁶I use SDL’s naïve modal semantics for simplicity. SDL faces a number of challenges for modeling natural language modals (Kratzer 1977, 1981), but these challenges generally concern embeddings that are irrelevant for present purposes.

⁷ R is serial iff $\forall w \in W . \exists v \in W : wRv$.

- If all maximally accurate acts at all possible states assign credence n to some proposition p , then *rationality* requires assigning credence n to p .

This generates our translation schema. Different interpretations of the TRUTH NORM determine different corresponding indomitable acts. For each interpretation I , look to the credence function Cr_v^I of your counterpart in each I -deontic alternative v to your world w . Having Cr_v^I is an indomitable act at w according to I .

Then we can ask: which Cr_v^I can genuinely play the roles of an indomitable act? If accuracy is understood as divergence from the truth: from *which* truths?

2 World-relative interpretation

A first-pass interpretation of the TRUTH NORM: you ought to believe p in w iff p is true at w . Call this the *world-relative interpretation* of the TRUTH NORM. It determines the following accessibility relation for ‘ought’:

WORLD-RELATIVE: $\forall w, v \in W, wRv$ iff for all $p \in \mathcal{P}$: $w \in p$ iff $v \in Bp$.

That is: v is a deontic alternative to w just in case in v you believe all and only truths about w .

The world-relative interpretation faces a problem. Because you’re not omniscient, there’s some true proposition p that you don’t believe. This constrains deontic alternatives:

in w	in any v s.t. wRv
p	you believe p
you don’t believe p	you believe that you don’t believe p

But this means that at all deontic alternatives to your world, you have a false belief. So you ought to have a false belief. But the TRUTH NORM says you ought to believe all and only truths. So the world-relative interpretation of the TRUTH NORM can’t be correct.

Sorensen (1988) calls true conjunctions of the form ‘ p and I don’t believe p ’ ‘blindspots’: they are possibly true propositions that aren’t truly believable (at least by agents who believe the conjuncts of conjunctions that they believe). Bykvist and Hattiangadi (2007) argue against the TRUTH NORM on these grounds: it violates an extremely weak ‘ought’-implies-‘can’ principle, according to which you ought to ϕ only if it’s metaphysically possible that you both ϕ and ought to ϕ .^{8,9}

⁸This objection is related to, but distinct from, the Church-Fitch paradox. Church (1945/2009) showed that on all normal modal logics, (ia) entails (ib), and the same form of argument shows that (iia) entails (iib).

- (i)
 - a. All truths are knowable.
 - b. All truths are known.
- (ii)
 - a. All truths are truly believable.
 - b. All truths are truly believed.

The Church-Fitch paradox generates problems for (ia) because (ib) is false. Does the analogous entailment cause problems for the TRUTH NORM? Unlike (ib), the TRUTH NORM is a *deontic* claim: even if (iib) is *actually* false, perhaps it *ought* to be true.

⁹This problem translates immediately to AFE. The world-relative indomitable act at w is having Cr^{WR} :

3 Absolute interpretation

The problem with the world-relative interpretation is that it assumes that in any ideal world v , you believe all and only truths about some other world: the nonideal reference world w . But if in v you believe all and only truths about w , then you aren't omniscient with respect to v .

The intuition behind the TRUTH NORM might simply be that, ideally, you'd be omniscient. So any deontic alternative v to w should be one in which you believe about all facts in v (rather than facts in w). To capture this interpretation of the TRUTH NORM, we need a different accessibility relation. Call this the *absolute interpretation*:

$$\text{ABSOLUTE: } \forall w, v \in W, wRv \text{ iff } \forall p \in \mathcal{P}, v \in p \text{ iff } v \in Bp$$

In English: the deontic alternatives to any w are the worlds v in which you're omniscient (i.e., in v , you believe all and only truths of v). This counts as *absolute*, rather than world-relative, because all worlds will share the same deontic alternatives.

On this accessibility relation, if v is accessible from w , then v is also accessible from itself. In other words, this accessibility relation is *shift-reflexive*. On all shift-reflexive frames, the following schema is valid: $\Box(\Box\phi \rightarrow \phi)$. (In English: it's obligatory that obligations be satisfied.) In ideal worlds, one satisfies one's obligations; otherwise, ideal worlds would find themselves wanting. (If even v finds itself wanting, then why would it be ideal to w ?) So w regards v as ideal only if v regards itself as ideal.

Notice that the world-relative interpretation doesn't satisfy shift-reflexivity. This was the heart of its problem: it entailed that an objectively epistemically nonideal world w will only regard another world v as ideal if you have false beliefs in v , and therefore v won't regard itself as ideal. So on that interpretation, the TRUTH NORM requires you to violate the TRUTH NORM.

The absolute interpretation nevertheless faces its own problems. First: the deontic alternatives to your (nonideal) world are worlds in which you're omniscient; and therefore, in those worlds, you correctly believe that you're omniscient. So at your world, you *ought* to believe that you're omniscient. So on the absolute interpretation, you're required to believe something false. This already runs counter to the informal, intuitive motivations for the TRUTH NORM. It's also on its face implausible that you, in your current situation, should believe that you're omniscient.

Moreover, insofar as you know introspectively that you suspend judgment on some propositions, on this interpretation, the TRUTH NORM requires you to believe something you *know*—by mere introspection!—to be false.

On the two interpretations we've considered, either you're required to have a false belief (on the world-relative interpretation) or you're required to believe something that's false (on the absolute interpretation). Formally:

$$C_r^{WR}(p) := \begin{cases} 1, & \text{if } p \text{ is true at } w \\ 0, & \text{otherwise} \end{cases}$$

But since you're non-omniscient, there's some true p you don't assign credence 1 to. So for you, having C_r^{WR} means assigning credence 1 to p and assigning credence 1 to the proposition *I don't assign credence 1 to p* . These credences are necessarily self-falsifying, so they can hardly be epistemically ideal. If the value that epistemic decision theories aim to maximize is accuracy in one's credences, that value can't be characterized in terms of proximity to self-falsifying credences. So having C_r^{WR} shouldn't be an indomitable act.

- *World-relative*: $\exists p \in \mathcal{P} \Box (Bp \wedge \neg p)$
- *Absolute*: $\exists p \in \mathcal{P} (\Box Bp \wedge \neg p)$

Each of these conflicts with the intuitions motivating the TRUTH NORM.

This objection to the absolute interpretation isn't airtight. One might respond: the absolute interpretation predicts that you objectively ought to believe you're omniscient—but that's because you objectively ought to *be* omniscient! If you were doing as you objectively ought, then you would both be omniscient and recognize your omniscience.

Still, a second, more serious problem arises for the absolute interpretation. Consider some contingently true proposition, e.g., that there are cats. One might expect that, on the TRUTH NORM, you ought to believe this proposition. But on this accessibility relation, this isn't the case!

Why? The absolute interpretation says that *any* world in which you're omniscient is a deontic alternative to the actual world. Not all of these will be cat-worlds. (Cats presumably aren't a necessary condition on your omniscience.) So at some deontic alternatives to w , you don't believe that there are cats. So, on our modal semantics, it's not the case that you ought to believe that there are cats.

The set of worlds where you're omniscient is orthogonal to most ordinary, contingent, first-order propositions. On the absolute interpretation, the TRUTH NORM doesn't require you to believe these propositions. It only requires that you believe propositions entailed by your omniscience.

In short: the absolute interpretation is inconsistent with the TRUTH NORM, in both letter and spirit. On the absolute interpretation, there are some falsehoods that you ought to believe, and it's not the case that you ought to believe most actually true, contingent propositions.¹⁰

4 Mixed interpretations

4.1 Mixed interpretation #1

The second, more serious problem for the absolute interpretation resulted from the assumption that the deontic alternatives to w are *all* worlds where you're omniscient. That includes worlds where the contingent facts are very different from the actual world. So it includes worlds where your omniscience requires belief in many propositions that are, in actuality, false. What if we restrict the set of deontic alternatives to avoid this worry?

A first attempt: perhaps the deontic alternatives to w are worlds where you're omniscient but the rest of the contingent facts are the same as at w .

¹⁰Again, these problems translate to AFE. Given the translation schema in §1.4, we look to the credence function of your counterpart in deontic alternatives to your world. Since, on the absolute interpretation, this includes a large range of worlds, with different characteristic probability functions, there is no unique, precise credence function that determinately counts as *the* indomitable act for your world, relative to which your epistemic ideality is measured. Still, the first problem translates over immediately: all candidates assign credence 1 to the proposition that you're omniscient. For the second problem, if there's no unique, precise, determinate indomitable act, then there's no unique, precise, determinate standard relative to which your epistemic ideality is measured. So on the absolute interpretation, for propositions that aren't entailed by your omniscience, AFE's normal story about how to measure accuracy doesn't apply. It may be that your accuracy is imprecise, or indeterminate—even relative to a specific accuracy measure.

MIXED INTERPRETATION #1: For all $w, v \in W$, wRv iff

(i) for all $p \in \mathcal{P}$, $v \in p$ iff $v \in Bp$, and

(ii) for all $p \in \mathcal{P}$ such that $p \not\equiv \text{you're not omniscient}$, $v \in p$ iff $w \in p$.

I call this a ‘mixed interpretation’ because the accessibility relation restricts the deontic alternatives suggested by the *absolute* interpretation in a *world-relative* way.¹¹

Clause (i) ensures your omniscience in any deontic alternative v to w , while clause (ii) ensures that v and w are alike in all facts compatible with your omniscience: for example, there are cats in v iff there are cats in w . The proposition *there are cats* doesn’t entail *you’re not omniscient*, and is true at the actual world. So by clause (ii), it’s true at all deontic alternatives to the actual world. So by clause (i), you truly believe it at all deontic alternatives to the actual world. And so you objectively ought to believe that there are cats.

Mixed interpretation #1 looks promising. It still predicts that you ought to believe you’re omniscient—which may or may not be a problem (see §3)—but it’s able to predict that you ought to believe many contingently true propositions. So it apparently improves over both the world-relative and absolute interpretations.

But this interpretation still yields unacceptable verdicts. Consider (1):

(1) You’re omniscient or Hillary Clinton is the US president.

(1) doesn’t entail your non-omniscience, and (1) is not true at the actual world. So, by clause (ii), at any deontic alternative v to the actual world, (1) is false. But clause (i) ensures that any deontic alternative v to any world is a world where you’re omniscient. So at every deontic alternative, (1) is true. So there are *no deontic alternatives* to the actual world.¹² The point here is quite general. So long as \mathcal{P} includes a first-order proposition that is both false at w and compatible with you being omniscient, there are no doxastic alternatives to w .¹³

One might respond: if the actual world has no deontic alternatives, that can hardly be a problem for the TRUTH NORM. The absence of deontic alternatives entails the (trivial) truth of the TRUTH NORM.

But it also entails the trivial truth of the FALSEHOOD NORM: believe all and only falsehoods. So the trivial truth of the TRUTH NORM should be no consolation to its proponents. More generally: it would be *logically impossible* to do as one ought. Nothing would be permissible.¹⁴

4.2 Mixed interpretation #2

Mixed interpretation #1 is too strong, and has a natural weakening:

¹¹In AFE, this means that the mixed, indomitable act for w is one which assigns credence 1 to all p s.t. $w \in p$ and $p \not\equiv \text{you're not omniscient}$, and is transparent to itself (i.e., $Cr^{M_1}(\text{your credence in } p \text{ is } n) = 1$ iff $Cr^{M_1}(p) = n$.)

¹²For AFE: w has no indomitable act.

¹³Given the assumptions in §1.2 about \mathcal{P} , it must include the union of the relevant proposition with the proposition that you’re first-order omniscient, which is all the example requires.

¹⁴It therefore violates the seriality condition that SDL places on deontic accessibility relations.

MIXED INTERPRETATION #2: For all $w, v \in W$, wRv iff

- (i) for all $p \in \mathcal{P}$, $v \in p$ iff $v \in Bp$, and
- (ii) for all $p \in \mathcal{P}$ s.t. $p \neq$ you're not omniscient, if $w \in p$, then $v \in p$.

Mixed interpretation #1 uses a biconditional in its second clause, while mixed interpretation #2 uses a conditional.

Given this accessibility relation, from the fact that (1) is not true in our world, we can't infer that it's not true in the deontic alternatives to our world. So this interpretation avoids the problem we posed for its predecessor.

But this variant faces its own problems. Many worlds—including our world—still have no deontic alternatives. To see this, let \mathbf{n} be the sequence of actually winning Powerball numbers tomorrow. Here are three propositions true of our world:

- (2) \mathbf{n} will be the winning Powerball numbers tomorrow.
- (3) If I believe that \mathbf{n} will be the winning Powerball numbers tomorrow, then I'll buy a Powerball ticket today with \mathbf{n} .
- (4) I won't buy a Powerball ticket today.

None of these entails that I'm not omniscient; so all are true at any deontic alternative to our world. But they jointly entail that I fail the omniscience condition at any deontic alternative to our world. ((3) and (4) entail that I don't believe (2), but (2) is true. So I fail to believe something true.) Clause (i) requires any deontic alternative to be one in which I'm omniscient. So our second mixed interpretation also predicts that our world has no deontic alternatives.¹⁵

4.3 Mixed interpretation #3

We can avoid the problem for mixed interpretations #1 and #2 if we can guarantee a nonempty set of deontic alternatives for any world. Let O be the set of worlds where you're omniscient. The selection function $\sigma : W \rightarrow O$ maps each world to the uniquely closest (i.e., most similar) world where you're omniscient.¹⁶ This ensures that the deontic alternative to w is very much like w : if there are cats in w , then there are cats in $\sigma(w)$.

MIXED INTERPRETATION #3: $\forall w, v \in W$, wRv iff $\sigma(w) = v$.

On Stalnaker's (1968) counterfactual semantics, Mixed Interpretation #3 claims: what you *ought* to believe is what you *would* believe, were you omniscient.

Will this do the job? What might a deontic alternative to our world be like? It depends on the similarity relation over worlds. Determining an appropriate similarity relation over worlds (for some purpose) is often a theoretical minefield. (See, e.g., Lewis 1979.) We'll consider various options.

¹⁵Or again: our world has no indomitable act.

¹⁶This rules out ties, which would generate further problems for this interpretation. For propositions about which tied worlds differ, this interpretation can't issue requirements.

4.3.1 Future similarity, future omniscience

Suppose the most similar *O*-world to ours preserves similarity to our future, except that you become (and remain) omniscient. That world, *v*, is a world where you hide your extra knowledge. The future would be profoundly different if you told the world how to cure cancer or halt climate change!

This means that *v* is a world where you're weirdly secretive. You don't tell your mother the winning Powerball numbers; you don't share the cure for cancer. Here are some other truths in *v* which, since you're omniscient in *v*, you believe of yourself in *v*:

- (5) I'm weirdly secretive.
- (6) I'm ungenerous with my mother.
- (7) I allow people to needlessly die of cancer.

So, on this interpretation, in the actual world, you objectively ought to believe these (presumably false) propositions. That's absurd. And, again, this requirement seems to run afoul of the traditional motivations for the TRUTH NORM.

All versions of mixed interpretation #3 face another objection: that in the deontic alternative to the actual world, (8) is true, hence believed, hence something you actually ought to believe of yourself:

- (8) Most of my empirical beliefs are entirely unsupported by my evidence.¹⁷

4.3.2 Present similarity, future omniscience

Let's try an alternative similarity relation, holding fixed more of the present but allowing the future to diverge. In particular, we'll hold fixed your actual willingness to share your beliefs, even if that means that the future of the actual world's deontic alternative will be quite different (cf. Lewis 1979). Then you needn't believe (5)–(7), because those propositions wouldn't be true in *v*.

This generates its own problems. Here are some other truths in *v* which, since you're omniscient in *v*, you believe of yourself in *v*:

- (9) I'll tell my mother the winning Powerball numbers.
- (10) I'll disseminate the cure for cancer.

So, on this interpretation, in the actual world, you objectively ought to believe these (presumably false) propositions. That's absurd.

4.3.3 Present similarity, future non-omniscience

In the previous two similarity relations, you elected not to act on your omniscience, or you acted on it. Your omniscience infected facts outside of your beliefs, facts that the TRUTH NORM shouldn't require or even apply to. The resulting variants of the TRUTH NORM entailed that you should believe some propositions that are actually false *and* (unlike the proposition that you're omniscient) not

¹⁷Alternatively, for Williamsonians about evidence: *Most of my empirical beliefs have no basis in observation or testimony.*

required by the TRUTH NORM to be true. Can we construct a similarity relation that quarantines the TRUTH NORM away from such propositions?

Let's pay attention to time-indexing: you ought_t to ϕ just in case you would ϕ at t if you were omniscient at t . We might use a similarity relation where you're omniscient *only at t*, after which your body of beliefs contracts back to your actual belief state. You therefore have no time to act on your omniscience in ways that affect the future: you have no chance of disseminating the cure for cancer. So you're not responsible for needless cancer deaths: before you can act, t is passed and you no longer remember the cure for cancer. So you're not required to believe either (7) or (10); similarly for the other examples.¹⁸

But this tweak generates its own problems. Here is a truth in v which, since you're omniscient in v , you believe of yourself in v :

(11) I'll lose almost all of my beliefs within the next minute.

So, on this interpretation, in the actual world, you objectively ought to believe this (presumably false) proposition. That's absurd.

4.4 Lessons

An important lesson: it's impossible to cleanly separate out propositions that are related to your beliefs from those that aren't. There is *no bright line between first-order and higher-order beliefs*. Propositions (5)–(11) are partly higher-order: whether you're secretive, ungenerous, etc. depends on what you believe and how you act on your beliefs.¹⁹ They aren't entailed by, or immediately about, your omniscience. But were you omniscient, your omniscience would make them true (on one or another plausible similarity relation), which would affect your beliefs about them. So each counterfactual interpretation of the TRUTH NORM predicts that you actually ought to believe some of these propositions, even though they're false. This can hardly capture the idea that you ought to believe all and only truths.

This objection differs from the worry that you ought to believe:

(12) You're omniscient.

This is false (alas!), but if the TRUTH NORM is correct, it ought to be true. The only reason why (12) is false in the actual world is that you're not living up to your objective epistemic obligations.

The same can't be said of (5)–(11). The TRUTH NORM doesn't require you to be weirdly secretive, or to tell your mother the winning Powerball numbers, or to lose almost all of your knowledge in the next minute. The falsity of (5)–(11) is no epistemic shortcoming of yours. So on these interpretations, the TRUTH NORM requires you to believe false propositions that the norm doesn't require to be true.

¹⁸Thanks to Kevin Dorst for suggesting this hypothesis to me, and for discussion.

¹⁹This is a different point from the claim defended in Hedden and Dorst (2021), who argue that almost all evidence is higher-order evidence. Their view isn't that a surprising amount of evidence has higher-order *contents*; rather, it's that a surprising amount of evidence is evidentially relevant to *other* attitudes with higher-order contents (in the context of their paper, contents about evidential support relations). Their view therefore depends on some controversial assumptions about evidential support, normative uncertainty, and evidence (non)transparency that are orthogonal to the arguments in this paper.

We've canvassed several interpretations of the TRUTH NORM; each yielded unacceptable verdicts. Interpreting the TRUTH NORM should be starting to seem like a degenerating research program. I inductively conclude: there's no credible interpretation of the TRUTH NORM. The TRUTH NORM is false.

5 Approaching truth

5.1 Accuracy maximization

Accuracy-first epistemologists can respond: couldn't there be a plausible alternative to the TRUTH NORM for the objective epistemic 'ought'? Maybe the TRUTH NORM went wrong by demanding perfect omniscience. Instead, we might take the objective epistemic 'ought' to demand *maximizing* your accuracy. Compare with expected utility theory: there, the subjective 'ought' prescribes maximizing expected utility, and the objective 'ought' prescribes maximizing utility *tout court*.

Similarly, AFE evaluates doxastic states for *degrees* of accuracy. Where omniscience is logically impossible, you objectively ought to maximize your accuracy to the extent possible—even if in doing so, your beliefs aren't *perfectly* accurate.

ACCURACY NORM: you objectively ought to have the most accurate possible beliefs.

This is a constitutive commitment of consequentialist accuracy-first epistemology (e.g. Greaves 2013; Caie 2013; Pettigrew 2016).

Like the TRUTH NORM, the ACCURACY NORM has world-relative, absolute, and mixed interpretations. The ACCURACY NORM pairs best with the *world-relative* interpretation. On the absolute and mixed interpretations (but not the world-relative interpretation), deontic alternatives would still be worlds in which you're omniscient, and so these interpretations would face precisely the same objections as with the TRUTH NORM. The ACCURACY NORM is interesting only to the extent that it's distinct from, and avoids the problems of, the world-relative TRUTH NORM.

A temptingly simple formulation, which Castro and Vassend (2017) appear to endorse:

ACCURACY NORM (WORLD-RELATIVE): In w , you ought to have the most accurate doxastic state compatible with being in w .

Problem: this norm is only nontrivial if we use "small worlds" that don't fully specify the agent's doxastic states. And where these small worlds specify some of the agent's doxastic states, the norm sometimes produces unacceptable consequences. Consider a toy example: you only have attitudes towards these two propositions and their negations:

- p
- q =: I don't believe p .

We'll use a tripartite belief model involving belief, suspension of judgment, and rejection. I assume that belief in truths and rejection of falsehoods are both more accurate than suspension of judgment

on either truths or falsehoods, which is in turn more accurate than belief in falsehoods and rejection of truths.

Let w_{pq} be a world where both p and q are true. At w_{pq} , ACCURACY requires you to perform this epistemic act:

- believe q ; reject $\neg q$; and
- suspend judgment on p ; reject $\neg p$

Why not believe all truths and reject all falsehoods at w_{pq} ? Believing p is incompatible with being in w_{pq} : if you believe p , q is false, and hence you're not in w_{pq} . It's logically impossible to truly believe both propositions that are true at w_{pq} . Given that you can't have the perfectly accurate attitude toward p , suspension is the best you can do. The rest of your attitudes can then be perfectly accurate.

Problem #1. These beliefs are incoherent. Coherence requires that if you suspend judgment on p , you also suspend on $\neg p$. But if you suspend judgment about both p and $\neg p$, then your credences do not maximize accuracy: they are less accurate than the incoherent combination. Hence, the world-relative interpretation of the ACCURACY NORM requires having incoherent doxastic attitudes.²⁰

Problem #2. Your options for the maximally accurate beliefs to have at w_{pq} were limited by the fact that we specified some facts about your beliefs in our identification of w_{pq} . This generalizes: the more fully described worlds are within your space of propositions, the more constrained the options will be for maximizing accuracy. If you can form beliefs about your attitude toward any proposition p , the maximally accurate attitude toward p at any world w will be whatever attitude you happen to have at w . If worlds are specific enough to determine your total belief state, then the most accurate belief state compatible with w is *whatever your belief state in w happens to be*. (This is, of course, the *only* belief state compatible with w !) So the relevant accessibility relation is identity:

ACCURACY (WORLD-RELATIVE) for non-toy examples: $\forall w, v \in W, wRv$ iff $v = w$.

It would therefore be *impossible* to violate the ACCURACY NORM—rendering the norm trivial.

5.2 State-relative accuracy-maximization

We need a different characterization of what beliefs are 'possible' for the ACCURACY NORM. Accuracy-first consequentialists understand available epistemic acts as decision theoretic options. To model these, instead of using coarsened-grained "small" worlds, we use both fine-grained *worlds* and coarser-grained *states* of the world, which form a partition S over W . Each state of the world is compatible with your holding different beliefs.²¹

²⁰See Caie 2013 and Castro and Vassend 2017 for a defense; Castro and Vassend use a credal analog of this example.

²¹This characterization is consistent with different views about how to construct the relevant partition: states could entail all and only propositions that aren't about one's beliefs, or they could entail *some* propositions about one's beliefs, or they could fail to entail some propositions that aren't about one's beliefs. My objection will extend to all of these views.

	s_1	s_2	\dots
adopt belief state 1	w_1^1	w_1^2	\dots
adopt belief state 2	w_2^1	w_2^2	\dots
\vdots	\vdots	\vdots	\ddots

Instead of a world-relative interpretation of the ACCURACY NORM, we use a *state-relative* interpretation:

ACCURACY NORM (STATE-RELATIVE): For any $s \in S$, if you're in s , you ought to have the most accurate doxastic state compatible with being in s .

STATE-RELATIVE: For all $w, v \in W$, wRv iff

- (i) $\exists s \in S$ s.t. $w, v \in s$, and
- (ii) $\forall p \in \mathcal{P}$, if $s \subseteq p$, $v \in Bp$.

In English: v is a deontic alternative to w iff v is in the same state as w and in v , you're omniscient about all propositions entailed by that state.

Notice: whatever facts related to belief states are entailed by s will constrain what beliefs are maximally accurate at s . For example, if s entails p and that no one believes p , the most accurate total belief state relative to s will demand suspension of judgment on p . How then is this an improvement over the world-relative interpretation? Since the partition of states can be coarser-grained, states need not entail higher-order contents (i.e., propositions about your beliefs). Your actual beliefs need not be the most accurate beliefs compatible with the state. So satisfying the norm is nontrivial.

Problem: the only way to block the result that we always trivially have maximally accurate beliefs is to use states that are silent about some of our doxastic attitudes. This means that there are propositions that we are perfectly capable of entertaining, but maximizing accuracy doesn't require us to be correct about them.

What objectively ought we believe about our own beliefs? Even if it were possible to sharply distinguish first- and higher-order beliefs, it's implausible that higher-order beliefs are *entirely unconstrained* by accuracy. So should we have states entail propositions about first-order beliefs? Second-order beliefs? First-order beliefs about cats, but third-order beliefs about watermelons? This interpretation of the ACCURACY NORM is partition-dependent, and any choice of partition will be unacceptably arbitrary.

5.3 Ramifications for AFE

These problems for the ACCURACY NORM infect any form of AFE. Still further problems afflict versions of AFE committed to:

PERFECTIONISM: Our doxastic states are better to the extent that they approximate what we objectively ought to believe.

PERFECTIONISM is widely assumed in epistemic decision theory.²² If we ought to have the maximally accurate doxastic state in a world or state, then:

Problem #1: This entails that we are sometimes, perhaps always, objectively *and* subjectively required to be incoherent. Consider a credal variant of the example above: one has attitudes toward p and its negation, and toward q^* =: *No one has credence in p greater than .5*, and its negation. At w_{pq^*} , the maximally accurate credence function is c :

$$c(p) = .5; c(\neg p) = 0; c(q^*) = 1; c(\neg q^*) = 0$$

If we do better at w_{pq^*} by approximating c , then as Castro and Vassend (2017) show, c^* is *accuracy-dominated*:²³

$$c^*(p) = .5; c^*(\neg p) = .5; c^*(q^*) = 1; c^*(\neg q^*) = 0$$

Castro and Vassend conclude that c^* is always irrational, evidence be damned. Not only will incoherence be objectively required, as in §5.1, but it will also be subjectively, *rationally* required. Surely, though, there are evidential situations in which you have good reason to have credence .5 in both a proposition and its negation, and confident that everyone you're quantifying over has no higher credence. (For example, every ordinary toss of a fair coin!)

Problem #2: Sometimes distinct belief states will be tied for most accurate at a world or state. How do we determine which we should use to measure the inaccuracy of other beliefs? These choices will be arbitrary, and their arbitrariness will infect verdicts about rationality.

The ramifications for AFE discussed here focus on problems associated with identifying what ideal to approximate at each world. But AFE faces other problems associated with treating the doxastic states under evaluation as facts within worlds. As Carr (2017, 2019) argues, these facts can be causally, probabilistically, and constitutively relevant to other facts in the world, and epistemic decision theories that are sensitive to these forms of act-state dependence won't uphold traditional AFE results: arguments for probabilism, conditionalization, the principal principle, and even strict propriety break down.²⁴

6 There is no objective epistemic 'ought'

We've looked hard, but found no plausible way of understanding the TRUTH NORM. Each candidate interpretation placed absurd requirements on believers. Similarly for the ACCURACY NORM. So I inductively conclude: there is *no truth or accuracy norm* on belief.

The TRUTH and ACCURACY NORMS seem to be the weakest viable candidates for a general objective epistemic 'ought'. (Analogous knowledge-based norms would face similar problems.) If even these have counterexamples, then I conclude: there is *no objective epistemic 'ought'*.

One might object: aren't there still weaker candidates for objective epistemic 'ought's? Let's consider a few candidates.

²²I take the name from Pettigrew (2016); his (2018) is, I believe, committed to the claim as stated.

²³Castro and Vassend's argument is meant to refute Joyce's (1998; 2009) arguments for probabilism. The dispute may be reconstructed as a disagreement about which interpretation of the objective epistemic 'ought' to feed into PERFECTIONISM.

²⁴See also Briggs (2009) for an earlier, related discussion of Dutch book arguments for putative norms governing higher-order states, including strong forms of Reflection and introspective omniscience.

6.1 Avoid Falsity; Know that which You Believe

Here are some candidates for weaker objective epistemic norms:

AVOID FALSITY: You ought not have false beliefs.²⁵

KNOW THAT WHICH YOU BELIEVE. You ought to know that which you believe.²⁶

Unlike the TRUTH NORM, these norms don't require you to believe anything: *ipso facto*, they don't require you to believe that you're omniscient, responsible for cancer deaths, or any of the other implausible prescriptions of the TRUTH NORM.

But these norms still produce implausible verdicts. First, on both norms, at every deontic alternative, you have no false beliefs; so at every deontic alternative you do not falsely believe that you have any false beliefs. So both norms entail that it's objectively impermissible to believe that you have any false beliefs. For similar reasons, KNOW THAT WHICH YOU BELIEVE also entails that it's impermissible to believe that you have any unjustified or Gettiered beliefs.

If you're like me, you have powerful inductive evidence that you do have false, unjustified, or Gettiered beliefs (though you don't know which). If you do, on these proposals, it's impermissible to recognize your own fallibility. So, there are *truths* that you're prohibited from believing. Again, though, this might not be a problem. The proposition that you're fallible is true only because you violate one of these norms.

More problematically: both norms are so weak that they can't play various roles that one might want an objective epistemic 'ought' to play, as set out in section 1.1: what we ought subjectively to approximate in the pursuit of epistemic value, or what it would take to realize the greatest possible epistemic value. AVOID FALSITY only requires avoiding false beliefs; KNOW THAT WHICH YOU BELIEVE only requires not having beliefs that don't count as knowledge. One could satisfy both norms without having any true beliefs, by suspending judgment on everything.

Suppose one of these two norms is such that satisfying it is sufficient for doing as you objectively epistemically ought. This is inconsistent with two claims that many contemporary epistemologists accept—especially those who have been tempted by the TRUTH NORM:

- (i) You objectively epistemically ought to maximize epistemic value.
- (ii) True belief has more epistemic value than suspension of judgment.

Why? Given (ii), if you suspend judgment on everything, you don't maximize epistemic value: having some true beliefs, without trading in any false beliefs, will have higher epistemic value. So, given (i), suspending judgment on everything is insufficient for doing as you objectively epistemically ought.

Worse: universal suspension—global skepticism—is guaranteed to *perfectly* satisfy AVOID FALSITY and KNOW THAT WHICH YOU BELIEVE. Forming any beliefs risks violating these norms. So a form of weak dominance argument shows that, on each of these norms, forming any beliefs whatsoever is *irrational*. You're *subjectively* epistemically required to be a skeptic!

²⁵Defended by Raleigh (2013). Thanks to Alan Hájek for an equivalent suggestion.

²⁶This norm is discussed by Williamson (forthcoming); Littlejohn (forthcoming); Simion et al. (2016). Thanks to an anonymous referee for suggesting discussion of it.

The upshot: AVOID FALSITY and KNOW THAT WHICH YOU BELIEVE don't have the right shape to play the role of an objective epistemic 'ought'.²⁷

6.2 *Become omniscient*

One might think that the problem lay in our starting assumption: that the TRUTH NORM evaluates beliefs you already possess. Rather than saying that you ought to already *be* omniscient, perhaps the correct truth norm says that you ought to *become* omniscient. This will avoid many of the original TRUTH NORM's problems. After all, it doesn't require you right now to believe any falsehoods ('I will disseminate the cure for cancer'). If you *become* omniscient, then the relevant propositions won't be falsehoods.

BECOME OMNISCIENT NORM: what you ought_t to believe_{t+} is what you would believe_{t+} if you were omniscient at t^+ (where t^+ is, say, a moment after t).

What does this requirement amount to? In the actual world, I'm sad to report, you will not be omniscient a moment from now. Does this mean that you will violate your epistemic obligations? It's unclear. BECOME OMNISCIENT imposes present obligations on your future attitudes: you're obligated now, at t , to become omniscient in the future, at t^+ . So you're not *currently* violating your current (t) obligations. Since you won't be omniscient at t^+ , you will violate your t obligations. But you won't violate your obligations at t^+ , because at t^+ BECOME OMNISCIENT only imposes requirements on your behavior at a future moment, t^{++} . So if BECOME OMNISCIENT is the strongest objective epistemic norm, then bizarrely, there is never any fact of the matter about whether you're doing as you ought!²⁸

6.3 *Alternative semantics for the Truth Norm*

All interpretations of the TRUTH NORM we've considered were modeled in terms of possible worlds semantics, the orthodox form of modal semantics for modal logics and, in generalized form, natural language semantics. What if the problem isn't with the TRUTH NORM, but with possible worlds semantics?

²⁷In the sense set out in §1.1. The 'ought' of these norms is *externalist*, to be sure, but not *objective* in anything like the traditional sense. Note that many forms of normative language that are tempting to classify as 'objective' turn out to exhibit information-sensitivity; see Kolodny and MacFarlane (2010) and especially their unfortunately unpublished (Kolodny and MacFarlane, manuscript). Going beyond the agent's evidence isn't sufficient for an 'ought' to count as objective.

²⁸Contrast with ordinary future-directed norms: if your mother tells you at 6:12 that you must eat your broccoli, then there's some future time interval (possibly contextually determined)—say, 6:12 to 6:15—during which you can fulfill that obligation. But throughout that time interval the norm applies to you, not as a future-directed norm, but as a present-(interval)-directed norm. When 6:15 rolls around, you're in violation of your obligation. (For more on the quasi-referential properties of tense, see Partee (1973); Ogihara (1995).)

If the BECOME OMNISCIENT norm functioned in this ordinary way, then when t^+ rolled around, it would require you to believe at t^+ what you would believe at t^+ if you were omniscient—that is, it would partly function as a *present*-directed norm. And then, for its present-directed verdicts, it would suffer the same problems as mixed interpretation #3. In order for it to avoid those problems, it must be (somehow) *exclusively* future-directed—in which case it suffers the problems in the main text.

Consider an example of this style of proposal. One commitment of SDL, Kratzer’s modal semantics, and their relatives is that necessity modals like ‘ought’ satisfy Agglomeration:

Agglomeration: $\lceil \text{ought } \phi \rceil, \lceil \text{ought } \psi \rceil \vDash \lceil \text{ought } (\phi \text{ and } \psi) \rceil$

One might take the arguments above to support the idea that the TRUTH NORM is non-agglomerative. Then one can use a variant of the world-relative interpretation. For a true proposition p that you don’t believe, this variant predicts that you ought to believe p and you ought to believe that you don’t believe p , but it’s not the case that you ought to hold both beliefs. So it needn’t commit to the claim that you ought to have a false belief.

The problem with this proposal is that it cannot serve the purposes of the objective epistemic ‘ought’ that I set out in section 1.1. It won’t be able to characterize ideal total doxastic states; it won’t serve as a standard by which to measure the value of total doxastic states. And for these reasons it’s hard to see how it could inform a theory of the subjective epistemic ‘ought’: coherence norms apply at the level of total doxastic states. A single belief, in isolation from a backdrop of mental states, cannot be evaluated for coherence.

More generally: the problems I’ve raised don’t depend on possible worlds semantics as such. They concern very general questions: What specific beliefs are required by the TRUTH NORM? Are the contents of those beliefs actually true? What does conformity to the norm amount to? Can someone conform to the norm while having false beliefs or suspending judgment? I’m skeptical that any interpretation of the TRUTH NORM, in any semantic framework, can provide acceptable answers.

To understand the broader semantic problem: some form of ‘ought’-implies-‘can’ is widely accepted (even if the relevant ‘can’ is extremely weak). What ought to be the case is restricted by what can, or could, have been so, given some fixed set of circumstances that constrain what an agent ‘can’ do. What circumstances should we hold fixed for the truth norm? If we hold fixed all of the truths of the believer’s world, then we hold fixed truths about the believer’s ignorance. So the ideal beliefs will either truly or falsely self-ascribe ignorance about at least one proposition; in either case, the agent is non-omniscient. If we don’t fixed all truths of the believer’s world, then there will be some truths that they’re permitted—perhaps even required—not to believe. So ideal beliefs will sometimes be non-omniscient. Both options defeat the purpose of the TRUTH NORM.

7 The bigger picture

7.1 Ignore the higher-order

All of the problems I’ve raised for truth norms have concerned higher-order beliefs and credences. And so there’s a temptation to set those aside—to restrict the scope of the TRUTH NORM to the first-order. This is no solution.

First, it’s wildly ad hoc. Recall, many propositions that aren’t primarily about one’s beliefs are nevertheless species of higher-order belief contents: beliefs about one’s generosity, lottery prospects, responsibility for cancer deaths, etc. Cordoning off *all* beliefs with partly higher-order contents is both impracticable and unjustifiable. One would somehow have to justify claiming that the truth about propositions about whether Tom is red-haired objectively epistemically matter, but the truths about whether Tom is secretive don’t.

Second: truth norms are meant to serve certain purposes in epistemology, discussed in 1.1: for example, in AFE, characterizing ideal epistemic acts, where the epistemic value of other acts may be measured relative to this ideal. This in turn determines which epistemic acts are *rational*, relative to an epistemic decision theory. If a truth norm doesn't apply to beliefs with even partly higher-order contents, then these don't affect a belief state's accuracy or epistemic value. As a result, AFE will place no constraints—objective or subjective—on higher-order beliefs. It won't even prohibit *explicitly contradictory* higher-order beliefs! The idea that one could be perfectly rational while holding explicitly contradictory beliefs strains credulity.

For many purposes—e.g. philosophy of science—restricting our attention to the first-order is unproblematic. But for epistemology, and assessments of agents' overall epistemic rationality, such a restriction is indefensible. Epistemic rationality concerns *all* of our doxastic states, individually and as they interact with each other.

On the traditional picture of scientific inquiry, the scientist stands outside the system she investigates. When we broaden the scope of beliefs under evaluation, we are necessarily ourselves inside the system under investigation—our world. We seem to want to apply the TRUTH NORM to agents who are somehow 'outside of the world', so that conforming to its prescriptions wouldn't affect the world in ways that undermine the norm. I believe there's some kernel of truth in this—but it concerns the nature of evidence and evidential support relations, not ideal or rational beliefs.²⁹

7.2 Value without norms

If there's no TRUTH NORM on belief, does that mean that epistemic norms are entirely independent of truth-related considerations? Well, no. The falsity of the TRUTH NORM and its various weakenings doesn't give a reason to reject the idea that true belief is *valuable*. We can (partially) order worlds in terms of alethic value (value deriving from truth or accuracy). Worlds in which you are closer to omniscient are in important ways *epistemically better* than worlds where you have more uncertainty, which are in turn better than worlds where you have more mistaken beliefs. Alternatively—and to my mind, preferably—we can (partially) order doxastic states, construed as abstracta (say, credence functions, or 'belief functions' from propositions to B or $\neg B$) in terms of alethic value. For example, if a credence function c 's credence in p is closer to the truth of p than c^* 's for any $p \in \mathcal{P}$, then c has greater alethic value than c^* .

It's just that this value has no *deontic* import. There's no objective norm associated with this value. We can't say that you ought to maximize the value of your beliefs. For example: let p be a true proposition that you don't believe. At the actual world, then, the most alethically valuable doxastic state includes belief in both p and $q =: I \text{ don't believe } p$. But we can't conclude on this basis that you *ought* to have those beliefs. If you did, you'd be guaranteed to have a false belief.

Nevertheless, the comparative alethic value of worlds, or of doxastic states, may still constrain the *subjective* 'ought' of epistemic rationality. It may, for example, be that you subjectively ought not adopt a doxastic state that is accuracy dominated by another. And so even if there is no truth norm on belief, the norms of epistemic rationality may still be governed by truth-based considerations. If so, though, the way in which they are so governed will differ greatly from pragmatic decision theories and consequentialist theories of pragmatic rationality in general.³⁰

²⁹Details of this proposal are beyond the scope of this paper; see Carr (2017) and Carr (2019).

³⁰See Carr (2017) for elaboration.

How, then, should we think about the relation between objective and subjective norms? Objective *pragmatic* norms are generally understood in terms of omniscience: very roughly, what you objectively ought to do is what you subjectively ought to do if you were omniscient. But the objective ‘ought’ doesn’t extend from the pragmatic to the epistemic. While the pragmatic realm makes room for both subjective and objective ‘oughts’, the epistemic realm only has room for the former; and as I showed in §5, this has serious implications for the possibility of importing pragmatic norms (e.g. decision theories) into epistemology. We can use omniscience to structure objective prescriptions, but we can’t prescribe omniscience.

Is it correct to say that you should believe the truth, then? No.

References

- Alston, William. 1988. The deontological conception of epistemic justification. *Philosophical Perspectives* 2. 257–299. doi:10.2307/2214077.
- Briggs, R.A. 2009. Distorted reflection. *Philosophical Review* 118(1). 59–85.
- Briggs, R.A. and Richard Pettigrew. 2018. An accuracy dominance argument for conditionalization. *Noûs*.
- Bykvist, Krister and Anandi Hattiangadi. 2007. Does thought imply ought? *Analysis* 67(4). 277–285.
- Caie, Michael. 2013. Rational probabilistic incoherence. *Philosophical Review* 122(4). 527–575.
- Carr, Jennifer Rose. 2017. Epistemic utility theory and the aim of belief. *Philosophy and Phenomenological Research* 95(3). 511–534.
- Carr, Jennifer Rose. 2019. A modesty proposal. *Synthese* 198(4). 3581–3601. doi:10.1007/s11229-019-02301-x.
- Castro, Clinton and Olav Vassend. 2017. Ideal counterpart theorizing and the accuracy argument for probabilism. *Analysis* 78(2). 207–216.
- Chrisman, Matthew. 2008. Ought to believe. *Journal of Philosophy* 105(7). 346–370. doi:10.5840/jphil2008105736.
- Church, Alonzo. 1945/2009. Referee reports on fitch’s ‘a definition of value’. In Joe Salerno (ed.), *New essays on the knowability paradox*, 13–20. Oxford: Oxford University Press.
- Easwaran, Kenny. 2016. Dr. truthlove: or, how i learned to stop worrying and love bayesian probabilities. *Noûs* 50(4). 816–853.
- Greaves, Hilary. 2013. Epistemic decision theory. *Mind* 122(488). 915–952.
- Greaves, Hilary and David Wallace. 2006. Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind* 115(459). 607–632.
- Hedden, Brian and Kevin Dorst. 2021. (almost) all evidence is higher-order evidence. *Analysis* 82(3). 417–425. doi:10.1093/analys/anabo81.

- Ichikawa, Jonathan. 2024. *Epistemic courage*. Oxford.
- Ichikawa, Jonathan Jenkins. 2022. You ought to have known: Positive epistemic norms in a knowledge-first framework. *Synthese* 200(5). 1–23. doi:10.1007/s11229-022-03872-y.
- Joyce, James M. 1998. A nonpragmatic vindication of probabilism. *Philosophy of Science* 65(4). 575–603.
- Joyce, James M. 2009. Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of belief*, vol. 342, 263–297. Dordrecht: Springer.
- Kolodny, Niko and John MacFarlane. 2010. Ifs and oughts. *Journal of Philosophy* 107(3). 115–143.
- Kolodny, Niko and John MacFarlane. manuscript. Ought: Between objective and subjective.
- Kratzer, Angelika. 1977. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy* 1(3). 337–355. doi:10.1007/BF00353453.
- Kratzer, Angelika. 1981. The notional category of modality. In H.-J. Eikmeyer and H. Rieser (eds.), *Words, worlds, and contexts: New approaches in world semantics*, Berlin: de Gruyter.
- Leitgeb, Hannes and Richard Pettigrew. 2010a. An objective justification of bayesianism i: Measuring inaccuracy. *Philosophy of Science* 77(2). 201–235.
- Leitgeb, Hannes and Richard Pettigrew. 2010b. An objective justification of bayesianism ii: The consequences of minimizing inaccuracy. *Philosophy of Science* 77(2). 236–272.
- Lewis, David K. 1979. Counterfactual dependence and time’s arrow. *Noûs* 13(4). 455–476.
- Littlejohn, Clayton. forthcoming. A plea for epistemic excuses. In Julien Dutant and Fabian Dorsch (eds.), *The new evil demon problem*, Oxford University Press.
- Nelson, Mark T. 2010. We have no positive epistemic duties. *Mind* 119(473). 83–102. doi:10.1093/mind/fzp148.
- Ogihara, Toshiyuki. 1995. The semantics of tense in embedded clauses. *Linguistic Inquiry* 26(4). 663–679.
- Partee, Barbara H. 1973. Some structural analogies between tenses and pronouns in english. *The Journal of Philosophy* 70(18). 601–609.
- Pettigrew, Richard. 2016. *Accuracy and the laws of credence*. Oxford University Press UK.
- Pettigrew, Richard. 2018. Making things right: the true consequences of decision theory in epistemology. In H. Kristoffer & Jeffrey Dunn Ahlstrom-Vij (ed.), *Epistemic consequentialism*, 220–240. Oxford: Oxford University Press.
- Raleigh, Thomas. 2013. Belief norms and blindspots. *The Southern Journal of Philosophy* 51(2). 243–269.

- Rosenkrantz, Roger D. 1981. *Foundations and applications of inductive probability*. Atascadero, CA: Ridgeview Press.
- Simion, Mona. 2023. Resistance to evidence and the duty to believe. *Philosophy and Phenomenological Research* 108(1). 203–216. doi:10.1111/phpr.12964.
- Simion, Mona, Christoph Kelp and Harmen Ghijsen. 2016. Norms of belief. *Philosophical Issues* 26(1). 374–392. doi:10.1111/phis.12077.
- Sorensen, Roy A. 1988. *Blindspots*. Oxford University Press.
- Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in logical theory (American Philosophical Quarterly Monographs 2)*, 98–112. Oxford: Blackwell.
- Williamson, Timothy. forthcoming. Justifications, excuses, and sceptical scenarios. In Julien Dutant (ed.), *The new evil demon: New essays on knowledge, justification and rationality*, Oxford University Press.
- Wrenn, Chase B. 2007. Why there are no epistemic duties. *Dialogue: The Canadian Philosophical Review* 46(1). 115–136. doi:10.1017/soo1221730000158x.