

Replies to Egan, Schroeder, and Harman

John MacFarlane

Author Meets Critics on *Assessment Sensitivity*, Central APA, 2016

Contents

Introduction	1
Egan	2
The big picture	2
Communication	3
Rejection vs retraction	5
Postsemantics	6
Schroeder	7
Is relativism too flexible?	7
On asserting future contingents	9
On relativism vs expressivism	10
Harman	11
Why not moral relativism?	11
What about moral uncertainty?	13

Introduction

My book has two interconnected aims. The first is to give a clear account of what it is to be a relativist about truth. Philosophers writing about truth relativism have tended to focus on refutations of the thesis or defenses against these refutations, with neither side saying very clearly what the thesis amounts to. I propose a way of understanding the thesis that makes it

philosophically interesting, but not so interesting that it becomes indefensible. I try to say clearly exactly what one needs to show in order to establish that truth is relative in this sense.

The second aim of the book is to make available a new way of understanding perspectival or subjective aspects of our thought and talk. There is a standard menu of options here, including forms of objectivism, contextualism, and expressivism. To these options I want to add another: an expression can be assessment-sensitive. That is, its extension can be sensitive, not to the context in which it is used, but to the context from which it is assessed.

These two aims are connected. I argue that relativism about truth is best understood as a commitment to the possibility of assessment sensitivity. Making sense of this commitment requires substantial philosophical work. The idea that a single assertion might be as true, as assessed from one context, and false, as assessed from another, is not one that we have a clear antecedent grip on. Accordingly, the book falls into two main parts. The first part (Chapters 1–6) is devoted to articulating and making philosophical sense of assessment sensitivity, the second (Chapters 7–12) to arguing that a number of philosophically interesting kinds of discourse exhibit assessment sensitivity.

The first part is *a priori* and conceptual. It tries to say what a linguistic practice that involves assessment sensitivity would look like. By doing this it makes coherent sense of relativism about truth.

The second part looks carefully at how we use language and argues that the best descriptive semantic theory for the target expressions takes them to be assessment-sensitive. I look at claims of taste, future contingents, knowledge attributions, epistemic and deontic modals, and indicative conditionals. In each case I suggest that the availability of an assessment-sensitive semantics helps defuse intractable debates between rival positions, each of which gets *something* right. In the final chapter, I consider why it might make some sense for these expressions to be assessment-sensitive.

Egan

The big picture

One small note on Andy's excellent presentation of the big picture.

I don't think norms for assertion and retraction are the *only* things we want from a theory of meaning. In general, we want norms governing speech acts that can be made with these sentences. I give an example focusing on assertion, but ideally we'd want to say something about questions, imperatives, &c. I don't think this way of thinking of the project is idiosyncratic. I argue in the book that it's pretty close to a conception one can find in David Lewis.

Andy says: "What the whole apparatus is in the business of doing is correctly characterizing the norms governing uses of sentences." Yes – this is a *descriptive* project. The innovation in the descriptive apparatus is important because, if we don't have adequate descriptive tools, we'll have an inadequate menu of options for how the language could work. This leads to intractable philosophical debates, e.g. between contextualists and invariantists, that might be resolved if we had assessment sensitivity.

Communication

Andy says: "phenomena about communication are very much not on the ground floor." That's correct: the word "communication" does not appear in the 8-page long index.

It's not that I don't think explaining communication is important, or that I don't concede that communicating is one of our main purposes in making assertions. It's just that I think these effects should be explained in terms of something more fundamental. I am drawn to think this partly because, although many of our assertions have a communicative intent, not all of them do. Assertions can be used to make jokes, take positions, bully people, demonstrate one's knowledge to a teacher, get the audience to see a connection, and so on. My thought is that we ought to explain their use in the transmission of information—a use that only some assertions have—in terms of some features that *all* assertions have. (Indeed, it's pretty weird for a philosopher to fixate on transmission of information as the central function of assertion. How many of the assertions made at this conference can be described as having transmission of information as their main goal?)

Andy's central question is this: "Why put assertion and retraction norms on the ground floor, rather than putting communicative role on the ground floor? (So why have the output of the pragmatics be speaker-side norms for

assertion and retraction, rather than uptake-side norms for acceptance or context update?)”

So, why not try to characterize the meaning of ‘tasty’ by fixing the update, instead of fixing assertion and retraction norms? Andy suggests this could be done as follows: the update rule for ‘vegemite is tasty’ is to self-attribute ‘being disposed to enjoy vegemite.’ (If you don’t want to perform this update, you’ll reject the assertion.)

Andy cites these advantages:

- A. This gives a simple account of communicative effects.
- B. It gives a clear and straightforward account of disagreement.
- C. It allows denial and assent to track acceptance and rejection of the central communicated content.

I’ll concede A, but I have some reservations about B and C.

Ad B:

We need to explain the kind of disagreement that takes place when I say “vegemite is tasty” and you says “it is not tasty.” Andy’s diagnosis is very straightforward: we are proposing to update the context in incompatible ways. And I agree that this is “a recognizable kind of disagreement.” But I wonder whether it’s the kind we want here.

We might think of supposing as a speech act whose intended update is to temporarily add a proposition to the common ground for purposes of exploring its consequences. (The “current state” is saved, so we can return the common ground to it after doing this.) At least this strikes me as what an update theorist should say about supposing.

Now, if Andy says “Suppose that Trump wins the nomination,” and I say, “No—suppose first that Trump doesn’t win the nomination,” then we’re proposing incompatible updates. And there is a kind of disagreement here, just as there is when I propose that we go to the movies while you would rather we go to the ballet. But it seems to me that there’s a more robust kind of disagreement in the tasty case. (Chapter 6 of my book is devoted to taxonomizing varieties of disagreement in a fine-grained way.) As I put it in the book: in the debate about whether vegemite is tasty, one party’s assertion precludes the accuracy of the other’s. There is no question of preclusion of accuracy in the supposition case. Of course, maybe Andy is

right that there really is nothing more going on in the ‘tasty’ disagreement than in the supposition case, but I’m trying to get something more.

Ad C:

Andy objects that on my account, it makes sense to reject an assertion of “this is tasty” while simultaneously accepting “the central communicated content of the assertion.” The central communicated content, he says, would be (roughly) that the taste of the demonstrated food is pleasing to the speaker, since this is the condition that must obtain if the speaker is conforming to the truth norm for assertion.

I agree that we can extract *a* content by asking: what needs to be true if the speaker is conforming to the truth norm for assertion. But it seems misleading to call this the “central communicated content.” If my aim is to communicate to you that I like vegemite, I can say “I like vegemite.” In that case it makes sense to talk of the “central communicated content,” since the point of the assertion is to communicate a certain content—that the speaker likes vegemite. If instead I say “Vegemite is tasty,” then I might still convey that I like vegemite, but my assertion aims to do more than merely this. Since the central point isn’t merely communicative, why should we think we can identify something called the “central communicative content”? And why is *this* content more central than the content of the assertion itself—the assessment-sensitive proposition that vegemite is tasty? So this objection doesn’t much bother me, but I may not have seen its force.

Rejection vs retraction

I try to do everything in terms of speaker-centered norms for assertion and rejection. Andy asks why not norms for (audience-centered) rejection, instead of, or in addition to, retraction?

This is something I briefly consider. I think one could go that way, and in some ways it might be cleaner. But I was impressed by how far I could get in understanding disagreement, which seems to be an interpersonal thing, with just the speaker-side norms, so I went with that as it seemed more parsimonious.

I don’t have any interesting general commitment to speaker-side vs audience-side norms.

Postsemantics

Simon Goldstein's question is why the postsemantics is there, mediating between the semantics proper and the pragmatics. Is this a notational convenience or an indispensable moving part?

Somewhere in between. I completely agree that it's not indispensable; one could rewrite the pragmatics in terms of truth at a context and index, without ever defining "truth at a context of use and context of assessment."

However, doing this makes your pragmatics coupled with messy details of the semantics. I find that the interposition of postsemantics gives you an appealing kind of modularity, a separation of concerns. All the account of assertion needs to know is the truth-at-a-context-of-use-and-context-of-assessment profile of a proposition, and it can spit out norms for asserting it and retracting assertions of it.

Suppose you change your semantics proper—for example, after adding an operator you may find you need a new parameter of the index. In my system, you need only make a small change to the postsemantics, and you don't need to change the pragmatics at all. Without postsemantics, we'd have to change the pragmatics every time we had a structural change to the semantics.

In thinking this is desirable, perhaps I'm influenced by an analogy with software design. Typically programmers try to encapsulate distinct functions into "modules" with clearly defined "interfaces." Internal changes can be made to the modules, and as long as the interface stays the same, none of the programs that use these modules need to know. That's a desirable feature.

In addition, the logical notions (logical consequence and diagonal logical consequence) are defined in terms of truth at a context of use and context of assessment. Sure, we could give extensionally equivalent definitions directly in terms of truth at a context and index, but these would be messy and would have to change every time the details of the semantics changed. Also, since we define both the logical notions and the norms for speech acts in terms of truth at a context of use and context of assessment, these are kept in sync with each other, so there are interesting relations between logical consequence and norms for assertion and retraction.

Schroeder

Is relativism too flexible?

I've argued that we need some additional structure in semantic and pragmatic theories if we are to properly describe some linguistic practices. Without this additional structure, I've argued, we end up having to fit these practices to one or the other Procrustean bed, always with bad results.

Mark's central worry is that the extra structure I've introduced gives us too much flexibility to describe linguistic practices. For example, in my book I discuss a toy example, "noy", which is like "now" but it denotes the time of the context of assessment, not the time of the context of use. No natural language contains an expression that works anything like this.

Mark notes that nothing in my theory constrains which features of contexts of assessment can be semantically relevant. One could have a term whose extension is sensitive to the number of garbage cans present at the context of assessment.

I recall Jason Stanley making similar points when I first started talking about these ideas around 2003.

I think everything depends here on what we take the theory to be doing. I don't present semantic theories as explanatory theories. The point of giving a semantics, as I conceive it, is rather to *describe* a linguistic practice—a particular system of norms. And, given this conception of what I'm doing, I'm not sure there's any real problem with the flexibility to which Mark calls our attention.

Here's a kind of silly analogy. The semantic framework is like a typewriter. We can use it to write various sorts of stories. Now there's a story about Quine's typewriter (I think I got this from his NY Times obituary). The story is that he removed the question mark from the typewriter so he could add some logical symbols. "You don't miss the quantifier?" a reporter asked. And he replied: "Well, you see, I deal in certainties." Anyway, imagine a typewriter with no question mark key. You can write tons of stories on this typewriter. But there are some stories you can't write: stories with questions. (OK, not every question has a question mark, but most ordinary questions do.) Adding the question mark would allow us to write some stories that we couldn't write before.

Of course, adding the question mark also allows you to write out non-

sense that you couldn't write before, like '??H?e'. But it's hard to see this as a serious cost. After all, there were plenty of nonsense strings we could type out already, without the question mark. (Adding the question mark doesn't even change the cardinality of the set of nonsense strings we can type.) There might be a cost to adding the question mark key—if we don't want to remove an existing key, we need to make the keyboard bigger. But the mere fact that we can now type out nonsense stories that we couldn't type out before is not itself a cost. We wanted a tool for typing the stories we wanted to write. It's hardly a criticism of this tool that just mashing your hands randomly over the keys may produce nonsense.

I'm tempted to say the same thing about semantic theories. Adding a new dimension of contextual variability is like adding the question mark key to our typewriter. It allows us to describe linguistic practices we could not describe before. Some of these are practices that seem bizarre and are not realized in any natural language. They are like the nonsense strings we can type on the typewriter. But the fact that our addition allows us to describe bizarre practices that we could not describe before is not, itself, a cost. After all, *even before we added contexts of assessment*, our framework was already flexible enough to describe zillions of bizarre practices, even involving garbage cans. For example, nothing prevents us from describing a predicate that applies only to my left sock, the moon, and the number 7. I'm not aware of any natural language common noun that has an extension like this. Nothing stops us from describing a modal operator that requires its prejacent to be true at infinitely many accessible worlds. And nothing stops us from describing an indexical whose referent at a context of use is sensitive to the number of garbage cans present at that context.

So, one can't really reckon it a *cost* of the addition of contexts of assessment to our framework that it allows us to describe bizarre linguistic forms that we never see in natural languages. We could do that already, in spades, without contexts of assessment. And, if our aim is just to have a descriptive tool, it's not clear why it's a bad thing that we can do that.

Now, one might worry at this point that my conception of the project is too unambitious. Shouldn't semantic theory have an explanatory, and not merely a descriptive, goal? Shouldn't we aim to *explain* why, for example, natural languages contain certain kinds of modals, and not others?

Well, it would be great to explain this, and I'm not against trying to do it.

Maybe it is possible to get a set of semantic principles that constrain the space of possible meanings much more tightly than standard truth-conditional semantics, and do so in a way that is explanatory (in the same way that comparable syntactic theories are explanatory). I don't know how to do that, myself, and I'm not trying to do it. What's more, I don't see the enterprise I am engaged in—trying to find a framework rich enough to describe all of the meanings I think we actually find in natural languages—as incompatible with a more explanatory enterprise. Presumably it's a constraint on the explanatory enterprise that it explain why expressions have the meanings they actually do, so you might think a descriptive project is a necessary first propaedeutic to the explanatory project.

Mark's own suggestion, not surprisingly, is that expressivism is less flexible in a way that helps rule out garbage-can type meanings. On the relational expressivist interpretation, we can interpret "relativist" intensions—for example, sets of (world, taste, information state) triples—as telling us what would be consistent with the ordinary factual beliefs of a person who believed this proposition and had a certain taste and information state. This interpretation of the formalism, he suggests, constrains us to use parameters that correspond to psychological states believers can be in. So, garbage cans are ruled out.

Note, however, that if our aim is to explain why certain meanings are found in natural language expressions, and not others, we haven't gotten very far. Sure, we've ruled out one class of non-attested meanings (garbage-can contextual sensitivity). But we can still describe indefinitely many non-attested meanings. For example, one kind of psychological state someone can be in is hoping the Bears win. The relational expressivist framework allows us to describe expressions that are sensitive to this in ways we do not find attested in any natural language. (Of course, we have the predicate "hopes the Bears win," but this just expresses a descriptive belief about the subject.) So, I don't see that we've gotten very far.

On asserting future contingents

Mark thinks it's crazy that my norm for assertion, combined with my semantics, says that it's not strictly okay to assert future contingents.

I discuss this in Section 9.9. I don't think it's a terrible consequence. The

view doesn't prohibit us from using sentences that express future contingents from making assertions; it only prohibits us from asserting future contingent propositions. I can still say "I'll arrive on the 9:30 train," even if I think there's some small objective possibility of a train derailment or strike or earthquake, if the proposition I'm committing myself to is that I'll arrive on the 9:30 train, barring a train strike, earthquake, or other unpredictable but unlikely event. If you challenge me—"Even if there's an earthquake?"—I'll clarify what I meant, but normally there's no need to make all these pedantic qualifications explicit.

Even if asserting a future contingent proposition is always a mistake, we still need our semantics to tell us under what conditions such assertions must be retracted. What's distinctive about the relativist view is that it holds that, if you're at a future point where what was asserted came to pass, you don't need to retract your earlier assertion, even in the face of a proof that you didn't conform to the assertion norm in making it.

I also give another option, for those like Mark who find the view hard to swallow. The Indeterminate Relativist Postsemantics creates a gap in the normative predictions of the theory, allowing us to say that it's *indeterminate* (not settled by the theory) whether one may assert a future contingent proposition.

On relativism vs expressivism

At various points in the book I engage with expressivist views, acknowledging formal parallels with relativism and trying to get to the heart of the disagreement.

One thing I focus on as a key disagreement is this. An expressivist will typically say that to judge that *P* *just is* to be in a certain mental state: sometimes a belief state, sometimes a planning state, sometimes an intermediate credal state, depending on the content of *P*. So, for example, on Yalcin's account, to judge that it is as likely as not that it is raining is just to be in the state of having 0.5 credence that it is raining, while being sensitive to the question whether it is raining. On Gibbard's account, to judge that one ought to pack now is just to plan to pack now.

Relativist accounts do not make this identification. So, on a relativist account, the state of believing that there is as likely as not that it is raining is

distinct from the state of having 0.5 credence that it is raining, and the state of judging that one ought to pack now is distinct from the state of planning to pack.

Part of the motivation for expressivist accounts is that there should be a tight relation between these states. Gibbard complains that normative realists leave a gap between deciding what one ought to do and settling on what one is to do—as if you might conclude, okay, so the fact is I ought to pack. Now, what to do? But if the gap between judging that one ought to pack and planning to pack is too big for a certain kind of realist, it may be too small for the expressivist. Here I echo some worries of Tim Scanlon, who objected that Gibbard’s view makes it impossible to describe a certain kind of weakness of will.

The relativist view seems a happy medium. It establishes a tight normative connection between the belief and the non-belief state that the expressivist identifies. So it predicts that one is making a mistake if one believes that it is as likely as not that it is raining while not having an 0.5 credence that it is raining. But it does not *identify* these states.

So, that’s a pretty general difference between expressivist and relativist views. It will apply to any expressivist view that says: what it is to judge that *P* is to be in such and such a psychological state (other than full belief). Mark suggests that it’s not a deep objection, because one can construct better expressivist views to which this objection does not apply. I don’t understand all of what he says here (about credal and predal states). But it seems to me that an expressivist can only evade the objection by giving up the central idea of explaining what it is to judge that *P* as being in such-and-such a psychological state.

Harman

Liz Harman focuses on my treatment of “ought.”

Why not moral relativism?

“MacFarlane takes there to be a single moral truth.” (p. 3)

That sounds rather grand. The truth is less exciting. I simply don’t take a stand in the book on the question of “moral relativism.”

Liz is right that I don't make normative standards a contextual parameter. (I say "normative standards," not "moral standards," because my "ought" is an all-things-considered ought, not a narrowly moral one.) The semantics uses an information-relative ordering of choices, so changes in the contextually relevant information state can induce changes in the ordering. But what doesn't vary with context, on my account, is the way in which the ordering is determined on the basis of the information state. That's just fixed. But this isn't because I have knock-down arguments against moral relativism. It's because I was concerned to illuminate one dimension of the relativity of "ought"—relativity to information—and I wanted to keep things simple.

One could add an additional parameter, the "normative standard," and make the ordering relative to this in addition to the information state. And then one could debate about whether this parameter is "initialized" by facts about the context of use or the context of assessment. Saying it is initialized by the context of use would give us the kind of moral relativism Gil Harman has defended. I agree with Liz that it would be better to let it be initialized by the context of assessment. Although I don't myself argue for moral relativism, I think my framework gives moral relativists a better way to formulate their view, one that is more resistant to some of the usual objections.

Why do I not argue for moral relativism? I think that the sorts of arguments I use in the book to argue for an assessment-sensitive treatment of the thin evaluative term "tasty" don't carry over to terms of moral evaluation. In the case of "tasty," the objectivist position is difficult to defend, because claims of taste seem to be obviously subjective in some way. To a first approximation, we think we can call something tasty when we like its taste—even when we're fully aware that others have different tastes. So with "tasty," the problem is really one of deciding how to best characterize this subjectivity, and we have to choose between forms of contextualism, expressivism, and relativism. In the case of terms like "virtuous" or "morally right," by contrast, the objectivist position is much more defensible.

That's not to say there couldn't be a good argument for an assessment-sensitive treatment of moral language. But it would be a more subtle argument. One approach that seems to me promising is an indirect one. "Tasty" seems synonymous with "tastes good," so, as I suggest in the book, one

might investigate whether the assessment sensitivity of “tasty” is due to “tastes” or to “good” or to some interaction between them. Here it will be important to investigate “looks good,” “smells good,” etc., and also “tastes salty,” “tastes like chicken,” and so on. If one determined that “good” is the culprit, one might have an argument for the assessment sensitivity of “good person,” “good action,” and the like. I think we’ve really just scratched the surface here. (One of my grad students is working on this now.)

What about moral uncertainty?

I find this issue very interesting, and I’m grateful to Liz for provoking me to think about it.

It wasn’t obvious to me at first what my view implies about the debate between Uncertainty and Actualism. But after thinking about it, I’ve come to the view that, insofar as my view can accommodate moral uncertainty at all, it is in the Uncertainty camp.

So, what is the difference between Actualism and Uncertainty? They are competing views about what one ought to do in light of one’s full information (including moral information). Actualism says that the moral part of this information is irrelevant; if two information states agree about the non-moral facts, then whatever one ought to do in light of the one, one ought to do in light of the other. Uncertainty denies this, allowing that the moral component of the information is relevant to determining what one ought to do in light of that information.

Note: “information” here could be read doxastically, so that your information is your credences, or epistemically, so that your information is your knowledge or evidence. Liz doesn’t think it matters for her argument which way we understand “information,” but in her remarks she seems to be thinking of the doxastic reading. When I talk of information states in the book, I always think of them epistemically. We’ll come back to this.

Now, although I argue that bare “ought”s are assessment-sensitive, my theory does allow “oughts” to be explicitly relativized to a particular body of information. (The “for all x knows” operator, introduced in Chapter 10, shifts the information state to the information possessed by x at the time and world of evaluation.) So the issue of assessment sensitivity can be put off the table; we’ll just consider what my view says about the explicitly

information relativized oughts that are at issue between the Actualist and the Uncertaintist.

The most pressing question is whether my view allows for “moral information” and “moral uncertainty” at all. Liz says that my “bodies of information do not themselves contain *purely moral claims*, and no purely moral claims vary between contexts of assessment.” If this is right, then the question at issue between Uncertaintism and Actualism can’t even be stated in my framework! For these views are competing views about the truth conditions of “oughts” that are relativized to information states that include moral information.

So is it right to say that information states, in my view, can’t contain moral information? In a way, yes, and in a way, no. Let me explain.

We might think that what Fiona believes is this:

- (1) It is possible that I ought not have an abortion, and it is possible that I am permitted to have an abortion.

This is one way of representing her moral uncertainty. But in my framework, (1) is not going to be true. For it to be true relative to an information state i , that state would have to contain at least one world w_1 such that Fiona’s optimal choices relative to w_1 and i involve not having an abortion, and another world w_2 such that at least one of Fiona’s optimal choices relative to w_2 and i involve having an abortion. But remember, the ranking of choices in my semantics depends only on i , so the only way this can happen is if Fiona has different choices at w_1 and w_2 . If we assume it is settled by i what her choices are, then (1) cannot be true.

This might seem a bad consequence, and I discuss it extensively in section 10.7, “Modal ignorance.” I explain our temptation to say things like (1) as due to a fallacious pattern of reasoning with modals and conditionals: basically, we reason:

- (2)
 - a. if P , then she ought not to ϕ
 - b. if $\neg P$, then she is permitted to ϕ
 - c. but it is possible that P and possible that $\neg P$
 - d. so, it is possible that she ought not to ϕ , and it is possible that she is permitted to ϕ

This reasoning isn't valid, on the semantics I give for modals and conditionals, but it's easy to see why we might slide into this familiar pattern of reasoning. So this is an error theory explaining why we think that (1) could be true.

The upshot so far is that my system doesn't allow us to represent normative uncertainty in what might seem the most straightforward way. Deontic statements are global conditions on information states, so we can't represent uncertainty about a deontic statement.

But I don't think it follows that my framework is incapable of representing moral uncertainty. Consider what might naturally go in for P in (2a) and (2b), above. I would think it would be something like "abortion is morally wrong," "fetuses are persons," or something like that. These are not deontic claims, and there is actually nothing in my system that prevents us from thinking of them as propositions that are true or false at worlds. It's certainly not built into my conception of worlds that a description of a world fixes only the *physical* facts.

So, the way to get moral uncertainty into my picture is to allow the worlds to vary in their moral features as well as their physical features. (This requires a certain kind of moral realism, but I think that's a presupposition of the whole idea of "moral information," and thus of the Uncertainty/Actualism debate.) An information state that is uncertain about whether abortion is morally right will contain some worlds where abortion is morally right and others where it isn't. (Note that these worlds are going to be individuated epistemically, not metaphysically, so this does not entail that the moral status of abortion is metaphysically contingent.)

The question of Actualism vs Uncertainty now becomes the question whether the information-relative ranking ignores the moral component of information, in ranking choices, or takes it into account. If it ignored the moral information, the correct moral principles would have to be built into the ranking function itself (which is a fixed element of the semantics, not a parameter). This would give you a form of Actualism.

However, the form of Actualism you could get in this way, I now want to suggest, is very implausible. Here I need to raise a question for Liz, about the proper interpretation of Actualism. It seems to me that Actualism should be the view that whether "you ought to ϕ " is true depends on the moral facts at the world of evaluation (in Kaplan's sense, i.e. the world of

the index). Alternatively, you might say that it depends on the moral facts at the world of the context (the actual world). I think the former interpretation is much more plausible, since we want counterfactuals like

- (3) If abortion were wrong, then Fiona would not be permitted to have an abortion

to come out true. But the form of Actualism one could get in my framework is neither of these. It always uses the moral principles that are built into the view itself, rather than attending to the moral facts at the world of evaluation or the world of the context of use. I think this will lead to many bizarre predictions (to start with, the falsity of (3)).

A more plausible form of Actualism would require the ranking of choices to depend on the world of evaluation. But, recall, on my semantics, the ranking depends only on the information state. It doesn't know which world in the information state is the world of evaluation. So, if it is going to take moral considerations into account at all, it has to look at what is given by the information state, and this is just partial information, not the moral facts.

So I think I'm committed to Uncertainty. If it makes sense to talk of "moral information," then my view says it should be treated like any other information.

So, what about Liz's argument against Uncertainty?

I reject the key premise:

- (iii) S is blameworthy for ϕ 'ing only if S ought (in light of S 's information) to have refrained from ϕ 'ing

Here it matters how we think of "information." If S 's information = S 's credences (as Liz assumes in her discussion), then I reject this because I think people can be blamed for ϕ 'ing even when ϕ 'ing is what they ought to have done, in light of their credences.

Example. Suppose you have plenty of evidence that handling live rattlesnakes is dangerous. Ignoring this evidence, you have a high credence that handling the snakes is perfectly safe. So you handle the snakes—doing just what you ought to do in light of your credences. You're blameworthy for handling the snakes, I think. Yet, it's not the case that you ought, in light of your credences, have refrained from handling the snakes.

[In her paper “The Irrelevace of Moral Uncertainty” Liz makes a delicate distinction between being blameworthy for ϕ 'ing and being blameworthy for causing the ϕ 'ing. This is too delicate for me.]

Alternatively, we might try: S 's information = S 's evidence? (Liz says this works too.)

Then the problem is this. If your credences don't align with your evidence, you may act with evil intent but still end up doing what you ought to have done in light of your evidence. In this case you are blameworthy. Suppose your evidence makes it overwhelmingly likely that a vial contains necessary medicine, but you believe (against the evidence) that it contains poison. You give the vial to a sick person with the intention of killing her. Clearly, you are blameworthy for doing this. But it is not the case that you ought, in light of your evidence, to have refrained from giving the person the vial. Indeed, you ought to have given it.

In addition, on the epistemic reading of the subjective ought, I think the Uncertaintist could just reject (i). It's plausible enough that the slaveholder's credences support keeping slaves, but it's *much* less plausible that the slaveholder's evidence supports that. (What evidence do we have, that the slaveholder didn't, that shows slavery to be wrong?)