

Replies to Leite, Turri, and Gerken

Krista Lawlor

Stanford University

I am grateful to Adam Leite, John Turri, and Mikkel Gerken for their comments. J.L. Austin hoped that philosophy might be a collaborative exercise; in this symposium, collaboration takes the form of objections and replies. My commentators make widely different objections, so my replies are organized by author.

Reply to Leite

Adam Leite is perplexed about my approach to external world skepticism: why not say that we *know* that we're not the victims of global deception? The resources for saying we have anti-skeptical knowledge seem ready to hand in a relevant alternatives approach to knowledge. What deeper philosophical motives make me say we do not have such knowledge? And what distortion of my own framework has me saying this?

In reply, first, my motivation is simple: I take skeptical arguments from closure and from information-dependence seriously.¹ Admittedly Austin himself is not moved by skeptical argument—he sees it as rather ramshackle;² but then the key skeptical argument he confronted was the argument from illusion, which *was* ramshackle. Skeptical arguments from closure and information-dependence are challenging arguments, and not so easily dismissed. Crucially, they explicitly represent themselves as growing out of common epistemic judgments and practice. I think

¹ Closure based skeptical arguments are articulated by Dretske (1970) (1981); Nozick (1981); Stroud (1984); Vogel (1990), and form the basis of much contemporary epistemic discussion. See also Brueckner (1985); Cohen (1988) Hawthorne (2004). For the skeptical argument from information-dependence, see Wright (1985), (2002), who traces the idea to Wittgenstein.

² *Sense and Sensibilia* reveals a disdainful attitude to skeptical thought. “Ramshackle” is M.G.F. Martin’s phrase in his manuscript, “Austin: Sense and Sensibilia Revisited.”

that Austin would have found these arguments challenging, and worthy occasions for reflection about knowledge.³

You are enjoying a pleasant paddle on the lake. Looking across the water you see a Gadwall duck, floating peacefully. You reason that it's not a Wood Duck (it's too dull) or a Mallard (the bill is noticeably thinner). It is in fact a Gadwall, and your evidence eliminates all the reasonable alternatives, so you know it is a Gadwall. Now, do you also know it is not a very convincing decoy? Its not being a decoy follows from its being a Gadwall. (The first proposition implies the second.) Yet it would be natural for you to feel you need more evidence to know that it's not a decoy, natural to feel that the set of reasonable alternatives expands when knowing *this* is at issue (perhaps it's *not* a wooden decoy, but plastic yard decor?).

This natural pattern of epistemic judgment, and related pattern of linguistic behavior (claiming to know it's a Gadwall, hesitating to claim to know it's not a decoy) has been the focus of epistemology for some time now. Add the fact that it seems you should know what follows from what you already know, and we have all the ingredients we need for the closure based skeptical paradox: if you don't know it's not a decoy, then how can you know it's a Gadwall after all?

Leite's question is why not say that one knows that one is not the victim of global deception? My answer is that whatever we say here we should say about the more homely Gadwall case as well—the structure of the skeptical problems is the same. So would Leite have me hold that on looking across the lake your basic birding identification also enables you to *know* that it isn't a very convincing decoy? That doesn't square with common sense.

What I do say on behalf of common sense is this. In the ordinary run of things, you do not need to know that it is not a very convincing decoy when you know that it is a

³ In fact, I believe Austin (1946) was ahead of his time, anticipating something like the closure based skeptical dialectic when one considers his various remarks about the connection between knowing and "proving" that one is correct, within a relevant alternatives story about knowledge.

Gadwall. Common sense says we do not need to do the extra epistemic work it would take to know this further claim about decoys, in order to know it's a Gadwall. The explanation of how I manage to defend this common sense position takes more space than I have here. It involves showing how our commitment to closure is a commitment to a more nuanced principle that is sensitive to changed situations. It also involves Austin's insight that the truth of a knowledge claim depends on the situation it is made about, combined with a relevant alternatives account of knowledge. Crucially, *each target proposition brings its own alternatives to the table, and these alternatives may vary with the occasion*. So you can know it's a Gadwall, when it's being a decoy isn't part of the situation you are considering. That doesn't mean that on this same occasion you also know it is not a very convincing decoy. To know such a thing, you'd have to have more or different evidence to eliminate some new alternatives (it's being an animatronic robot, or a child's clever toy, or whatever alternatives would be reasonable in the circumstances).⁴

I take it that I am reclaiming what common sense would have us say about the homely skeptical Gadwall case. The same reasoning applies to more lurid cases involving handedness and radical skeptical hypotheses. In the emergency room after your chainsaw accident, the attending physician knows that you still have both your hands; being globally deceived (a brain in a vat, dreaming, what-have-you), is not a reasonable alternative in this situation. But this doesn't mean that the doctor knows that she is not so deceived.

In reply to Leite, then, there is no "distortion" of my thinking here. Other philosophers have invented other ways out of the skeptical paradox, none of which I find sufficiently commonsensical. That is why I advocate the Austinian resolution of

⁴ Contrary to Leite's recap of my position, I don't take a firm view of what *eliminating* an alternative is—I don't advocate the strict view of *ruling out* on which to rule out an alternative is to have evidence that entails its falsity, but suggest that a more liberal view is workable. Also, while I say that not-R is often a reasonable alternative when R is the target proposition for a knowledge claim, I don't mean it is the only alternative: often we will demand much more specific alternatives than this.

the paradox. The Austinian situation semantic theory combined with a relevant alternatives account of knowledge is *designed* to give this resolution of the paradox.⁵

At this point my question is, why is Leite so convinced that we should adopt a stronger stance toward external world skepticism? I believe I have retained all we want. We do not routinely have knowledge of anti-skeptical hypotheses, and we don't need it. Why should we go further and say that we know we are not globally deceived?⁶ Several possible answers are suggested in Leite's commentary. But I admit I am not entirely certain of what answer Leite means to give. Let me try to reconstruct a few possibilities.

(1) First, it could be that Leite assumes that *any* Relevant Alternatives Theory of knowledge must deliver the result that you know you are not globally deceived. This assumption can seem to be behind some of what he says. But the assumption is false. An RA Theory would deliver such a result through a series of claims that are entirely optional. For instance, one might claim that being demon deceived is a permanently irrelevant alternative for any proposition, and further claim that one knows the falsity of any irrelevant alternative.⁷ But these claims are rejectable by RA theories. (I think rightly rejected.) It is not an intrinsic feature of RA Theory that it must deliver knowledge of anti-skeptical hypotheses.

(2) Leite writes:

...there are two hands before me (and a computer and myriad other three-dimensional objects). That is decisive evidence that there is a world of things outside of my mind... And the crucial point here is that such considerations eliminate precisely the alternatives Lawlor's theory specifies as relevant to

⁵ I am perplexed when Leite says, "I will not challenge the basic structure of [Lawlor's] semantical proposal." If he understands my proposal, then he should not see my response to closure-based skepticism as a distortion of my thinking.

⁶ Note that there are pressures to say this, coming from Crispin Wright's discussion of skepticism. More about that soon.

⁷ Stine (1999)

the question of whether one knows that one is not a BIV, etc. [i.e. globally deceived].”

Leite claims that I know *that I am not globally deceived*, because the relevant alternatives to that proposition are eliminated. Why are they eliminated? It is not entirely clear.

(a) It could be Leite has a view here about *what the relevant alternatives are*. That is, whenever one considers the anti-skeptical target proposition *that I am not globally deceived* the only relevant alternatives are tame propositions such as *I am handless*. Such a tame alternative is easily eliminated (one might think), for instance when I look and see my hand. But why think that the only relevant alternatives to the proposition *that I am not globally deceived*, are propositions such as *I am handless*? Is it never reasonable to find other alternatives in need of elimination, in order to know that one is not globally deceived? I want to make room for the idea that given Descartes’ projects and interests, his knowing that his senses did not deceive him required ruling out some less tame alternatives. The possibility that he was merely dreaming everything, or that a powerful deceptive agent was feeding him these experiences, were relevant alternatives.

In some cases, perhaps it is true that the only relevant alternatives to a target skeptical proposition are quite tame. Leite himself works hard to create such a case—his Threatening Judge case *may* be one such (though I have serious doubts about this⁸). But it would be a significant mistake to suppose that once a situation arises in which a target proposition earns a set of relevant alternatives, those are

⁸ I have my doubts about the case, since a lot is going on in it, having to do with the practical exigencies of responding to an authority’s expressed wish for acquiescence (a Judge’s attempt at humor), so using it as data about routine knowledge claims is a bit tricky. Consequently, I hesitate to agree that “this example highlights, in many situations in ordinary life we quite reasonably expect, even demand, that others know that the standard skeptical hypotheses do not obtain. Leite adds: “In any ordinary context, would you really trust the reports of a person who claimed not to know whether she is awake or only dreaming?” But of course *claiming to not to know something* would raise separate issues of its own.

always the alternatives that will be relevant to the proposition on every further occasion. To illustrate this point, we need only consider cases where an expanded set of alternatives is newly made reasonable: you know that the trailhead is around the next curve—your map tells you that it is. Then your hiking buddy tells you the map has been revised. You know what time it is—you just looked at your watch. Then you read in the news that your brand of watch has issued a recall—is yours one of the series that was recalled? The moral here is a fact about relevance:

A target proposition does not have its relevant alternatives set permanently. Relevant alternatives to a target proposition will change from occasion to occasion.

Perhaps Leite misses this fact? His missing it would explain his taking his Threatening Judge to constitute a *prima facie* case, if not a “decisive proof” that we have anti-skeptical knowledge, period. He seems to think that by producing a case in which the relevant alternatives to a skeptical hypothesis (*that I am globally deceived, or a Brain-in-a-vat*) are quite tame, he has shown something of *prima facie* epistemic importance. It is hard to see what importance it has, once we recognize the aforementioned fact about relevance.⁹

(b) Let’s try again. As noted, Leite claims that one knows *that one is not globally deceived*, because the relevant alternatives to that proposition are eliminated. We

⁹ If Leite is correct, then when the judge demands, “I trust you know you are not dreaming!” and you are ready to avow, “I know I’m not dreaming, your honor”, the situation is one in which the judge deftly manages to hold fixed the set of alternatives to a limited “ordinary” sort (alternatives such as *that you are drifting off, that you have gone off your meds and are on the verge of unconsciousness (?)*). Even if we grant that one may truly say one knows *that one is not dreaming*, we must be careful of what epistemic morals to draw from this. No proposition is an island, Dretske says. Although we speak of knowing propositions, on the Austinian view, we might better speak of what we know in terms of *a proposition as opposed to a set of relevant alternatives*. So you know you’re not dreaming—you haven’t drifted off or gone off your meds. But do you know you’re not dreaming—you aren’t globally deceived? That would constitute a different epistemic achievement. Taking their cue from Dretske, Contrastivists labor this point. See Schaffer (2005). See also Yablo (2014).

are asking, why are they eliminated? We have set aside the possibility that they are eliminated because they are a fixed set of tame alternatives. So here's another answer. The relevant alternatives in question are eliminated because of the *evidential power* of (my knowledge) *that I have hands, that this is a computer, etc.* Leite says that a proposition such as *that there are two hands before me* is "decisive evidence" for the anti-skeptical proposition. By "decisive" I take it Leite means evidence sufficient for knowledge, evidence that eliminates all relevant alternatives to the proposition. (And given point (a) above, these relevant alternatives might include some non-tame alternatives.)

Why is the proposition *that there are two hands before me* decisive evidence, capable of eliminating the relevant alternatives to the anti-skeptical proposition?

Leite doesn't say. But we can try to reconstruct his thinking. Return to the Gadwall case with its skeptical hypothesis. You see the Gadwall across the lake, and know it's a Gadwall by its identifying marks. Knowing this, do you have "decisive evidence" that lets you know the anti-skeptical proposition *that it is not a decoy*? One reason to think that you do is this: *that it's a Gadwall* implies the falsity of any new alternatives that arguably become reasonable with the introduction of this anti-skeptical hypothesis—*that it is a Gadwall* implies *that it is not an animatronic robot, not a plastic toy* and so on. One might think, then, that this means you have good reason, or "decisive evidence" in favor of it's not being a decoy.

This is a tempting line of thought—Peter Klein for instance, defends it—but it won't do.¹⁰ Klein argues as follows: as a birdwatcher, you can know that it's a Gadwall on the basis of some evidence *e*, say, visual evidence from looking across the lake. The relevant alternatives to its being a Gadwall (it's a Mallard, a Wood Duck) are eliminated by *e*. Of course *e* all by itself does not eliminate skeptical alternatives, such as that it's an animatronic robot, or a clever plastic toy, etc. (Here Klein accepts the commonsense idea that *e* cannot *eliminate* an alternative if it is *consistent* with it: *e* is consistent with it's being an animatronic robot or a child's toy.) But, Klein

¹⁰ Klein (1995, 220)

thinks, although e doesn't eliminate these alternatives, once e gives you knowledge that it's a Gadwall, now you have a new reason for thinking the skeptical alternatives are false: *that it's a Gadwall* is reason to *think it is not an animatronic robot, or a child's toy, etc.*—reason enough to count as knowing that it is not, and so reason enough to count as also knowing that it is not a decoy.

The problem with Klein's vindication of anti-skeptical knowledge is that it imagines you might somehow squeeze more out of your initial evidence than is possible.¹¹ Your evidence e in favor of it's being a Gadwall doesn't let you eliminate that it is very compelling decoy, a life-like animatronic robot, a child's clever toy. But somehow, once you know it's a Gadwall, you now have decisive evidence against these skeptical alternatives. Magic!¹²

I think it would be good for Leite to say where he stands regarding the closure based skeptical argument. A wide-ranging discussion of the argument has been going on, and he seems to have opinions that engage with the theses under discussion. (Some of what he says is consistent with defending an unadorned closure principle for knowledge.) At the very least I think Leite's failure to engage with the paradox, and in particular with the response that contextualist relevant alternatives theory offers, is what makes it difficult for him to comprehend what I'm up to. Not engaging with a skeptical argument is a perhaps a philosophical position one can adopt, but it is harder to do so when one realizes that commonsensical judgments about evidence, knowledge, and relevance yield the resources for the argument.

Let me briefly turn to Wright's skeptical argument from information-dependence. Wright takes pains to show that ordinary patterns of reasoning are at work in the skeptical argument. He also begins with a homely case: you see a soccer ball kicked into the net. A goal! Do you thereby know a soccer game is being played? No—that's something you needed warrant for believing *before* your seeing the ball go into the

¹¹ Brueckner (2000, 146) writes, "...it is as if we are attempting to squeeze more out of S's evidence than is really there."

¹² Cohen (2000, 101) was tempted to call it magic: "At first blush, it looks as if Klein has pulled the rabbit out of the hat."

net can count as evidence that *a goal has been scored* (as opposed, say, to *a re-enactment of a famous goal having been executed*.) Wright next asks us to compare the reasoning in Soccer with Moore's famous proof of an external world:

Soccer

I Jones has just kicked the ball between the white posts

II Jones has just scored a goal.

Therefore

III A game of soccer is taking place.

Paraphrasing Wright, the argument is this:¹³ The tendency of I to support II is conditional on your possessing *independent* warrant for accepting III. It's only in the context of a genuine game of soccer that there is such a thing as scoring a soccer goal.¹⁴ Warrant for III doesn't *come from* II then, but partly constitutes the conditions under which II can itself be warranted by I. Wright asks us now to consider Moore's argument

I My experience is in all respects as of a hand held in front of me

II Here is a hand

III There is a material world (since a hand is a material object in space)

The skeptical complaint is that Moore's experience I is not an unconditional warrant, but provides warrant for II only on the condition that one already has collateral information, for instance, information that the experience recorded in I is unlikely to occur unless there is a material world. That is to say, one must have warrant for III

¹³ Wright (2002, 333).

¹⁴ Wright notes that if that context includes "collateral information" that scenes of sort I occur only when soccer games are actually being played (as opposed, say, to soccer films being rehearsed), then we have evidence for III from I. But in that case we would have evidence for III without going through II. Our collateral information would be sufficient.

already in hand. But apparently we don't have this warrant—how could we get it except by tracing the argument from I to II to III, which we now realize fails to transmit warrant, because circular?

Wright's skeptical argument is powerful because it generalizes to a wide variety of particular skeptical conclusions (about other minds, the past, and so on); but more importantly for our purposes, it is powerful because it is driven by common sense judgments about what constitutes cogent reasoning.

What I say in reply to Wright's skeptic, boiled down pretty savagely, is that that we do not always need positive information, or evidence, or reasons, in order to for I to warrant II. In some cases, the fact that there are no other reasonable alternatives to II, or that every other alternative to II is unreasonable, is enough for I to count in favor of II alone. We do not need warrant for III, then.¹⁵ Wright is correct that commonsense judgment won't let us argue from I to II to III. But we don't need to. We don't have warrant for III, but we don't need to in order for I to warrant II. (Here I accept that we could not get positive warrant for III without collecting evidence for it, of the kind provided by discrete encounters with the world provided by I and the inference to II. Leite does not seem to question this.¹⁶) The skeptical argument is defanged by showing that it is every bit as much commonsense judgment that one might be warranted in II on the basis of I, when there are no reasonable alternatives to II.

A reader may find my response to Wright's skeptic less than compelling. I cannot make my case fully here. What I hope to do here is to make clear that my motivation in addressing the argument is quite intelligible: Wright's skeptic presents a challenge that grows out of commonsense epistemic judgments about what constitutes cogent reasoning.

¹⁵ I am of course not the first person to pursue this line of thought. See Pryor (2004). My story about why one doesn't need warrant for type III propositions is different than Pryor's.

¹⁶ However, he might seem to question the idea that one needs to appeal to anything before proposition II. See below.

What would Leite have me say instead?¹⁷ Leite asks, why isn't it enough to "take at face value the sorts of things we count as evidence in ordinary life"? Well, on one understanding of this phrase, that is precisely what I do—you can take your experience of seeing your hand at face value and warrantably believe *that you have a hand*. But Leite wants more—he wants us to take II *that I have hands* as evidence that III *I am not the victim of global deception*. It's not clear how one could take II as evidence for III while also taking Wright's argument seriously. And again, I stress, Wright's argument is based on commonsense observations about cogent reasoning.

Leite seems to think we can know the anti-skeptical proposition III on the basis of II. But Wright's Soccer case makes it clear that this is a difficult position to uphold in general. Perhaps Leite thinks that Moore's argument is free of the flaw that Soccer exhibits? If that is the case, Leite has not told us why it is.

One general flaw that Leite complains about is that Wright's skeptic artificially "limits" our evidence and that "if we work with our ordinary notion of evidence—one on which considerations about the world often count as evidence—then it's simply false that the skeptical hypotheses "neutralize our evidence." But Leite doesn't make clear what he means when he says that skeptical hypotheses do not "neutralize our evidence." If he means that we don't ordinarily take skeptical hypotheses to be relevant alternatives for our ordinary perceptual claims, then of course I agree: we can know *that we have hands* on the basis of perceptual evidence as of hands. It seems he must mean that if one just lets oneself start with proposition II, *I have hands*, instead of with proposition I, *I have the experience as of hands*, then one can know that one is not the victim of global deception. But it is not clear how this is supposed to work. A parallel move in the Soccer argument would have us starting with seeing a goal scored and reasoning to the claim that a soccer game is being played. That bit of reasoning would leave your average soccer fan cold, rightly

¹⁷ Leite remarks that I note that Austin is concerned to reject the idea of an ultimate evidential basis of a uniform kind for all empirical claims. Why isn't Austin's rejection enough to dispense with Wright's skeptic? Because Wright's skeptic doesn't trade on this idea—he trades on the structure of cogent argumentation.

enough. It is one thing to support direct realism about perception, and another thing to suppose that direct realism constitutes an ambitious response to the skeptic—one that would give us knowledge of anti-skeptical propositions.¹⁸

Skeptical arguments from closure and from information-dependence are serious challenges. We can learn a lot about knowledge by taking these arguments seriously. Taking them seriously does not mean showing that we have anti-skeptical knowledge; nor does defending common sense require showing that we have anti-skeptical knowledge. Rather defending common sense requires showing that *we do not need anti-skeptical knowledge*. So I claim, anyway. This is no “distortion” of my thinking—it’s what my account is designed to show.

Reply to Turri

John Turri offers survey data in order to evaluate my account of assurance. He contends that the data suggest a negative evaluation, but I’ll focus on the question of whether the data bear on my account at all.

In experiments E1 and E2, Turri divides survey respondents into what he labels “Assurance” and “Assertion” conditions. He assumes (a) that if he deploys a prompt in which someone says, “I assure you” or “I know”, then the speech act in the scenario is interpreted by respondents as one of assurance. He also assumes (b) that if he deploys a prompt in which someone utters a declarative sentence and does not say, “I assure you” or “I know”, then the speech act is interpreted as an assertion.

Both these assumptions are incorrect. I explicitly point out that in our ordinary linguistic practice, assurances may be made with simple declarative utterances. I also point out that not every utterance of “I know” (and I might have added “I

¹⁸ Wright (ibid, p.345) himself responds to a similar objection made by John MacDowell, who argues that the I-II-III argument restricts evidence to claims made about subjective experience, while the direct realist will want evidence to include claims about the world; Wright responds that the failure of warrant transmission to support III afflicts us not because we restrict evidence to subjective experience, but because warrants are *conditional*.

assure”) is an instance of the speech act of assurance.¹⁹ These points are crucially important. It is very easy for consumers of survey prompts to understand an unadorned declarative utterance to be offering an assurance. If respondents in the “Assertion condition” are supposing that an assurance has been made, then it is not surprising that their mean response does not differ from those in the “Assurance condition.”

Furthermore, it is all too easy to imagine readings of particular prompts that render the responses irrelevant to disputes about norms of assertion and assurance. In the case of E1, it is easy to take the “assertion” or “assurance” as merely an effort to calm the listener, or to be reassuring. In the case of E2, it is easy to take the “see for yourself” as having nothing to do with checking up on the speaker; if Jeff does eat the leftovers he will see for himself, whether he was willing to accept Sally’s assertion/assurance or not. In the case of E3,

X can[not] blame Sally for being wrong

is ambiguous between

X can[not] blame Sally for Sally’s being wrong

which might be relevant if true, and

X can[not] blame Sally for X’s being wrong

which isn’t, and is a real distractor. Note that my claim is not that assurance gives take exactly the *same* kind of responsibility for *all* their hearer’s beliefs. Sally may have taken some kind of responsibility for the truth of her audience’s beliefs, but she

¹⁹ Here is the first footnote in the book: “One can of course say many things with ‘I know’, not all of them involving a claim to know. For instance, one can give encouragement to a listener, or indicate boredom (‘I know, I know...’)....” Turri excuses his use of the prompts saying, “[Lawlor] assumes that saying ‘I know’ a proposition counts as an assurance that the proposition is true.” Not only do I explicitly deny this assumption, the assumption is *bad*. And that’s what is really important here after all; if the assumption is bad, it ruins the experiment, whether I make it or not.

didn't in the same way take responsibility of the truth of her overhearer's beliefs. For instance X cannot sue for damages as an injured party, but he might still be in a position to enter the normative criticism that Sally didn't know and therefore *shouldn't* assure — anyone, and that's the point I am making here. It's not at all clear that this point is getting tested.

Turri's attempts to test whether his respondents in fact discriminated between speech acts, by seeing whether those given the "assure" language in their prompt would be more ready to say that the hearer in the prompt was "assured." It's pretty plausible that the result here is just a priming effect: people primed with a prompt that uses the phrase "assure" are more likely to select that term when subsequently prompted.

Given these failings the experiments lose interest. Nonetheless I must add a complaint about the way Turri misunderstands my text. The Key Test Statement in E1 and E2

[Listener] should recognize that, depending on how things go, [speaker] might change [speaker's] mind about p.

garbles what I say on p.14. Here's the crucial bit from my text:

When S assures that p [as opposed to in the case of mere assertion], it is not "recognized" [by S] that further evidence or circumstances may change S's opinion.

The context makes clear that I mean that an *assurer can't* "recognize" [openly acknowledge] *in advance* that further evidence or circumstances might lead him to retract p, since he takes himself to have suitably conclusive grounds for ruling out all the relevant alternatives to p's truth, and irrelevant alternatives aren't and shouldn't be "recognized" by him in this specific sense. The Key Test Statement replaces my claim about what the assurer can't recognize in advance (in the foregoing technical sense of "recognize"—note the scare quotes) with a claim about what the *listener should* recognize throughout: namely that things could in principle

unfold in such a way as to eventually force all concerned to “recognize” [come to regard] as possible something the assurer may have dismissed as irrelevant or disposed of on the basis of what he took to be conclusive grounds. That’s just fallibilism, and it’s entirely compatible with the relevant alternatives account, as I proceed to explain in the very paragraph in question.

Let me make two recommendations. First, one might confront a single test subject with a number of opportunities to give or withhold a particular description under different circumstances. Turri’s experiments confront each individual subject with only one scenario, in ways that make it impossible for the subject to guess what the questions are questions about. So the queries naturally attract irrelevant readings. Second: to test for the normative features and function of the speech act of assurance, we must be careful to have a scenario in which the function of assurance clearly *is* performed, and clearly is interpretable as such. In none of Turri’s scenarios is there any indication of concrete grounds the speaker might have for making the statement he or she makes, or any indication whether the listener accepts the assertion or assurance he or she is offered.²⁰ And it isn’t easy on any of the scenarios for listeners to quickly imagine grounds that would render the assertions or assurances in question acceptable. (There’s no suggestion that the speaker has advance knowledge of the size and shape of the detour in E1, no suggestion that the speaker remembers how old the leftovers are and might therefore have a warranted confidence that they’re still fine in E2, no suggestion that Sally could know about the prospects of NanoGenes in E3—is she an inside trader or ...?). So it’s easy for a respondent to feel that the assertion or assurance in question is unconvincing. At which point, it’s anyone’s guess what subsequent responses tell us about anything, let alone about the function of assurance and as opposed to assertion.

²⁰ So for instance, regarding E2, notice the naturalness of appending “but see for yourself” to an utterance that all can foresee does not perform the speech act of assurance: Cassandra assures Paris of his doom, knowing full well her remark will fail to perform the function of assurance. “...but see for yourself” she adds. The fact that one might add “but see for yourself” to an utterance that is not fulfilling the function of assurance, and isn’t expected to, is not relevant in testing the theory.

About the idea that knowledge is a norm of assertion: Defenders of the idea have had to make sense of the fact that we don't routinely blame people for assertions that are backed by something less than knowledge—justified belief will often get a person off the hook. Defenders of the knowledge norm have said that it is a norm the routine violation of which goes unremarked by ordinary folks. Bernard Williams tries to make sense of this claim. He finds it perplexing that there might be a norm that people regularly and blamelessly violate. Williams suggests that perhaps we can understand what makes it the case that the norm of knowledge applies in the case of assertion, despite its being regularly ignored by speakers, if we look to other speech acts that also constitutively require some kind of authority or warrant. He compares commands. People may try to give commands or orders without authority, and we can understand how they might seem blameless in doing so, if they think (incorrectly) that they have the relevant authority. So far, so good. But, Williams notes, explore the analogy a little further and things go wrong very quickly:

Commanding or ordering someone to do a certain thing is a special kind of telling-to, one done with authority as contrasted with mere ground-level telling-to, done without any special authority, as when an aggressive road-user tells another to get out of the way, or a helpful one tells another to look out. This ground-level telling-to parallels ground-level telling-that—merely telling someone that P, without any special entitlement. But ground-level telling-that is surely *already* assertion. (Williams, 2004, p. 77)

Here is how I understand Williams's argument (*Assurance*, p.51-52): it's only if the analogy holds between assertion and other speech acts that require authority, such as commands, that we will be able to understand how assertion might require authority—how it might be subject to a distinctive norm, despite the fact that this norm is routinely ignored. But the analogy breaks down, as soon as we consider it:

Command	Assertion
_____	-----
ground-level telling-to	?

For what is the ground-level, non-authoritative speech act that corresponds to the supposedly authoritative speech act of assertion? Any such ground level speech act will be 'already' an assertion—there is no speech act more basic, to which assertion is the authoritative counterpart. If Williams is correct, assertion does not require special authority, and does not require knowledge on the part of the asserter. In the book I recap Williams' argument, and leave it a good question for the defender of the knowledge norm.

What does Turri say in response to Williams' argument? In a footnote, he writes:

For the record, the argument that the analogy fails did not persuade me. For it overlooks a natural "ground-level telling-that" analogous to the "ground-level telling-to," which is supposedly the key missing ingredient that causes the analogy with commanding to fail. The natural analog is, quite simply, telling-to-believe. It's just a special case of a more general telling-to-do.

Turri's suggestion that assertion might stand to *ground-level telling-to-believe* as command stands to ground-level telling-to immediately encounters two objections that he doesn't bother to consider. First, there's no such thing as ground-level telling-to-believe. I can only be told to do things that can be voluntarily done, and believing isn't something I can do voluntarily. Second, even if there were such a thing it wouldn't solve the analogical equation. Assertion is not a kind of telling to believe, as it would have to be, if the analogical relation is to hold. Telling you to do a thing is simply futile if there's no hope of doing it ("I'm *telling* you to lift that building!"). But asserting something to you is not simply futile if I know you won't believe me. There can be lots of reasons to assert something to you even if I doubt it will produce any uptake in you. (As I write, I am painfully aware of some of them.)

Turri briefly suggests his own account of assurance as a "mode" of assertion. On Turri's account, assuring "... becomes sensible *only if* one's audience is apt to be

worried or anxious about things turning out a certain way.”²¹ It’s obviously a big task to defend one’s own account of assurance, and Turri didn’t have the time or space to articulate one, so I won’t discuss his claim. But I do notice that Turri’s own account will have difficulty with his own experimental data. Here is experiment E3’s prompt in its entirety: “Jeff and Sally are having a private conversation about which biotechnology firm is the best investment over the next year. Sally says to Jeff, ‘NanoGenes is the best investment, I assure you’.” Survey respondents apparently didn’t find it odd to say an assurance was given here, and it isn’t clear that Jeff has worries or anxiety. Perhaps Turri himself will want to consider the possibility that they didn’t take Sally’s utterance to be an assurance after all.²²

Reply to Gerken

Mikkel Gerken aims to argue against my theory. I am not certain what the argument is precisely, but this much is clear: he thinks that my theory makes two objectionable claims and that these claims are linked in such a way that if he calls one into question, that will call the other into question as well. The two claims are these:

(PF) Practical factors may (indirectly) affect the truth-value of knowledge claims

and

(Constitution) The constitutive function of knowledge ascriptions is to provide assurance.

I do advocate PF. What it takes to count as knowing may vary with the practical interests of conversation partners, because interests may impact what counts as relevant. Importantly, however, there are *limits* on how interests may impact what it

²¹ (My emphasis). Turri actually says “reassuring” here but I take it he hasn’t changed subject, and is still talking about assuring.

²² Thanks to David Hills for very helpful comments on this reply.

takes to count as knowing because *what is relevant is set by a reasonable person standard*. Additionally, the account is a species of intellectualist or purist theory, on which practical factors help to determine what is relevant, but the features of the subject that determine whether she counts as knowledgeable are all *epistemic*. To use a helpful analogy,²³ practical factors may affect whether I count as “fit” (by Navy Seal standards, not so much), while what determines whether I am fit are just my physical features (my blood pressure, resting heart rate, O₂ uptake and the like).

But I do not advocate Constitution. I think it’s pretty clearly false. In my first footnote I say,

One can of course say many things with ‘I know’, not all of them involving a claim to know [i.e. assurance]. For instance, one can give encouragement to a listener, or indicate boredom (‘I know, I know...’).

There are many utterances involving the use of the words “I know” that are not utterances where an assurance is given. Self-ascriptions of knowledge can figure in other speech acts than assurance. I note just a couple instances. There are also many utterances involving the use of the words “he/she knows” that are not cases where we communicate information about who has assurance-giving power.

So my reply to Gerken’s argument is that I reject the Constitution claim. It’s not clear who would ever accept this claim. And it’s somewhat perplexing that Gerken ascribes it to me.

Let me take this opportunity to do two things. First, it may help for me to briefly summarize my position about the connection between the semantics of knowledge ascription and the speech act of assurance. Second, it may help for me to briefly note some terminological choices that I made that could encourage misunderstanding. Finally, I’ll very briefly compare my project and Edward Craig’s project in *Knowledge and the State of Nature*.

²³ The analogy isn’t mine, and I regret that I can no longer identify where I first came across it in print.

(I) *The connection between the semantics of knowledge ascription and the speech act of assurance.*

The Austinian situation semantic theory is quite general in its application; in particular it is supposed to work for a wide variety of locutions, starting with simple declarative utterances. In Chapter 2 of *Assurance*, I use quantified expressions to illustrate how it works. It also works for knowledge ascriptions, which is what I go on to show. (Austin doesn't speak of knowledge ascriptions when he sketches his semantic observations, and I add some widgets to flesh out the story about knowledge ascriptions.) So, the first important point to note is (i) the Austinian semantic theory has general application to a wide variety of locutions.

The Austinian situation semantic theory may also be applied to knowledge ascriptions that are not made in the act of giving an assurance. For instance, it works for first-, second- and third-person ascriptions of knowledge that are not made in the service of assuring. So the second important point is (ii) the Austinian semantic theory has general application to all kinds of declarative utterances, *including non-assurance-giving ascriptions of knowledge.*

I do not in the book spend time giving examples or detailing how the semantics would apply in these other cases, assuming that it was enough to show how the semantics applies to *declarative claims* and to note that *some ascriptions of knowledge are merely declarative* (i.e. not all knowledge ascriptions give assurances). I did not imagine that readers would think that the semantic theory only applied to those knowledge ascriptions that figured in assurance-giving acts.

Note that I do *not* build the Austinian semantic theory for knowledge ascriptions by appealing to either the "constitutive" or non-constitutive functions of such ascriptions. I build it as Austin might have done, namely: first, by reading off the intuitively correct individual meanings of terms, and supposing that these term meanings compose in routine ways to determine the utterance's content (the "descriptive content" or "Russellian proposition"). I suggest that "knows ..." has an invariant meaning, roughly, "has conclusive reason to believe the true claim"

Second, I suggest that the situations that Austinian situation semantics should appeal to when it comes to utterances involving epistemic terms (like “knows”) have epistemic elements. Following Austin’s relevant alternatives theory of knowledge, I suggest that the situation one is talking about when one makes a knowledge claim is a situation that includes relevant alternatives.

I take it as a point in favor of a semantic theory that it comports well with the theory of speech acts. In particular, a semantic theory of knowledge ascriptions that comports well with a story about the speech act of assurance would have this as a point in its favor. While we have a choice of many different semantic theories of knowledge ascriptions, not all fare well in this regard. I try to show how a specific sort of *broadly contextualist relevant alternatives theory* comports well with our assurance-giving practice. This is what the Austinian account is—a broadly contextualist relevant alternatives theory.²⁴ It is a special sort of contextualism, namely situation semantics, and a special sort of relevant alternatives theory of knowledge, namely one on which a reasonable person standard determines what counts as relevant.

I imagined readers would wonder, *Why should we want yet another contextualist semantics? Or, Why should we take it that a relevant alternative is what a reasonable person standard determines to be relevant?* My answer is that the contextualist relevant alternatives theory I have to offer makes good sense of the data. And I take the data, broadly, to include both the usual sort of “variability” data that contextualists have long been rubbing our noses in, and a second sort of data, given much less attention, about our assurance-giving practice. What do I mean here, about the data provided by our assurance-giving practice? Here the thought is that assurance-giving has interesting and difficult to explain features—for instance: we give assurances even when we are in the dark about the particular needs of our interlocutors; we act as if we can give assurances even when we disagree about what matters; and we know the world can give us terrible turns, and make our

²⁴ For another instance of contextualism relevant alternatives theory, see Lewis (1996).

claims false, yet we give assurances happily over and over again despite this fact. My claim is that we can *explain* these features, and thus rationalize our practice, by thinking of knowledge claims as context sensitive (or better, “situation-sensitive”), and thinking of relevance as reasonableness. In this way, the specific theory I offer comports well with the facts about assurance-giving.

To recap: I nowhere say that all ascriptions of knowledge have a constitutive function, and the Austinian semantic theory does not appeal to constitutive functions of knowledge ascriptions. The semantic theory applies to all knowledge ascriptions. Neither I do motivate my contextualist relevant alternatives theory by appeal to the Constitution claim. Contextualist relevant alternatives theory is motivated by what has come to be called *variability data*. People are I hope further moved to adopt my account because it also offers a rationalizing explanation of our assurance-giving practice.

(II) Some terminological matters

I see that my preface may have encouraged a misreading. I’m grateful for the chance to clarify. There are a few specific confusions to clear up. Here’s what I say:

Wittgenstein and Austin both advise us to start our study of knowledge with knowledge claims, beginning with what one does when one says ‘I know...’. What is it that one does with such statements?

The simple answer is—one claims to know something.[footnote 1] But claiming to know is more than making a report about one’s epistemic position. In claiming to know, one offers one’s assurance to others. When one claims to know that one’s phone is working again after the outage, or that the meteor shower begins tonight, one offers an assurance to the effect that these things are true.

These first lines might encourage someone to think that whenever one says ‘I know’ one is doing the same thing, namely giving assurance. But I quickly clarify my meaning. The phrase “claiming to know,” is meant to be synonymous with “giving

assurance.” So the phrase “claiming to know” is not synonymous with “making a claim in which the word ‘know’ figures” or “self-ascribing knowledge.” I didn’t mean that every self-ascription of knowledge is an assurance.

As noted above, I further say in footnote 1:

One can of course say many things with ‘I know’, not all of them involving a claim to know [i.e. assurance]. For instance, one can give encouragement to a listener, or indicate boredom (‘I know, I know...’).

There are many utterances involving the use of the words “I know” that are not cases in which an assurance is given. There are many self-ascriptions of knowledge in which the act is not *claiming to know*.

These are quite a few terminological matters to throw at a reader in prefatory remarks. I was trying to avoid using the term “assurance” first thing, and focus readers on the idea that there is an action that correlates with some important uses of “know,” namely the act of claiming to know, or assuring. The basic idea is I hope clear: On the one hand, we have a broad class of *utterances in which the terms “I know” or “He knows”* are used to various effects. On the other hand, we have a special use of these terms, the special act of giving *assurances*. Call the former “knowledge ascriptions”, and the latter “assurances.” The act of assurance-giving *may* be done (often is done) explicitly by making a self-ascription of knowledge.

I go on to say that in the book my focus is on cases where assurance are given. This *might* invite the idea that when I go on in chapter 2 to give a semantics for knowledge ascriptions, it will only apply to those cases where assurance is given. But this would be a mistaken idea: the semantic theory of course applies more broadly than to just those cases where assurance is given. The semantic theory applies to knowledge ascriptions, whether or not they are made in the service of the act of giving an assurance.

Finally, introducing Austinian Semantics in chapter 2, I write:

Verbal assurances express claims to know. We have seen that an assurance is sound when the knowledge claim it expresses is true. Compare: your advice is sound when the wisdom it expresses is, in fact, a bit of wisdom. But when is a knowledge claim true?

When I say here that assurances express “claims to know” or (“knowledge claims”), I mean that assurances involve self-ascriptions of knowledge. That does not mean that the only speech act in which knowledge is self-ascribed is the speech act of assurance, of course. (I inadvertently here added a layer of complexity by using “claiming to know” for the *act* of assurance giving, and “knowledge claim” sometimes to mean *knowledge ascriptions*. Hindsight being perfect, I see this was unhelpful.)

For readers who are so patient as to still be reading, let me offer this simple restatement: we have on the one hand the speech act of assuring, and on the other hand, the self-ascription of knowledge, e.g. “I know”, which is often made in the course of performing the speech act of assuring. The semantic theory applies to ascriptions of knowledge quite generally, whether the ascription is made in the service of assuring or not.

(III) Appeals to function

Let me briefly add that my account can be compared to Edward Craig’s, though it differs in some key respects. Craig begins by asking, why do we have the concept *knowledge*, what function does it serve? His answer is that we have a need to identify reliable informants. The concept of *knowing* begins its life as the concept of *being an identifiably reliable informant*, Craig suggests, but he also suggests that concepts tend to undergo “objectivizing” pressure, under which they may lose some of their features or see them diminished in importance. Identifiability of the informant tends to be diminished under objectivizing pressures. He gives a sort of “state of nature” story about how the concept might have first arisen, and how our linguistic practices might support the role of the concept in fulfilling its function.

My account doesn't begin with the function of the concept of knowledge, but with the function of assurance-giving.²⁵ And it investigates those linguistic practices that support assurance-giving.

A more substantial difference between my view and Craig's is this. I do not suppose that we might directly derive the linguistic *meaning* of "knows" from the function of assurances. When I suggest an invariant meaning of "knows" and a relevant alternatives theory of knowledge, I do so for a variety of reasons: in part because these are Austin's suggestions and I'm interested in how we might build an Austinian view; in part because the suggestions have intuitive pull—both an invariant meaning of "knows" and a relevant alternatives theory of knowledge have been advocated by a number of theorists of very different persuasions, for good reason; finally, and most importantly, because doing so lets us build a unified theory of knowledge and knowledge ascriptions that makes sense of our assurance-giving practice. In this respect, then, concerning the meaning of "knows", or the content of the concept *knowledge*, Craig and I differ. Craig derives meaning directly from function; I do not.

There are broad similarities in our views. The central similarity of course is that I also think it important to keep function in view. The function of assurances, I claim, is to offer exclusionary reasons to the assurance's recipients, and it is a signal fact about self-ascriptions of knowledge that they figure in assurance-giving acts. Given some of the perplexing features of assurances, there are good questions about how

²⁵ I am less certain of the relation between concept and linguistic meaning, and would hesitate to equate them. About the linguistic meaning of "knows *P*" it seems a good hypothesis that it is roughly, "has conclusive reasons to support the true belief *P*." But the concept of knowledge may include more than this. (Gettier considerations might suggest as much.) If I were to hazard a view about the function of our concept of knowledge, I would say it includes inter-personal coordination. Craig's picture is of a person with needs for information, scanning the horizon for someone who can give him what he needs. There's something to this of course, but I think we are concerned with both sides of the exchange of sureties: assurance-seekers want not only to acquire a true belief, but to be relieved of concerns, and seek reasonable surety; assurance-givers are ready to take on responsibilities, giving surety, insofar as it is reasonable.

self-ascriptions of knowledge assist in the performance of assurance-giving acts. The answer to these questions requires a theory of knowledge and a semantic theory of knowledge ascriptions that makes sense of how assurances can perform their function. I think that every theory of knowledge and knowledge ascription has this explanatory burden.

Works Cited

- Austin, J. L. 1946. "Other Minds." *Supplement to the Proceedings of the Aristotelian Society* 20: 148–87.
- Brueckner, Anthony. 1985. "Skepticism and Epistemic Closure." *Philosophical Topics* 13: 89–117.
- . 2000. "Klein on Closure and Skepticism." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 98 (2): 139–51.
- Cohen, Stewart. 1988. "How to Be a Fallibilist." In *Philosophical Perspectives*, edited by Tomberlin, 2:581–605.
- . 2000. "Contextualism and Skepticism." In *Skepticism*, edited by Sosa and Villanueva. Boston: Blackwell.
- Dretske, Fred. 1970. "Epistemic Operators." *Journal of Philosophy* 67: 1007–23.
- . 1981. "The Pragmatic Dimension of Knowledge." *Philosophical Studies* 40 (3): 363–78. doi:10.1007/BF00646423.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Klein, Peter. 1995. "Skepticism and Closure: Why the Evil Genius Argument Fails." *Philosophical Topics*, Epistemology, 23 (1): 213–36.
- Lewis, David. 1996. "Elusive Knowledge." *Australian Journal of Philosophy* 74: 549–67.

- Nozick, Robert. 1981. "Knowledge and Skepticism." In *Philosophical Explanations*, 172–217. Cambridge, MA: Harvard University Press.
- Pryor, James. 2004. "What's Wrong with Moore's Argument?" *Philosophical Issues* 14 (Epistemology): 349–78.
- Schaffer, Jonathan. 2005. "Contrastive Knowledge." In *Oxford Studies in Epistemology*, edited by Gendler & Hawthorne, 235–71. Oxford: Oxford University Press.
- Stine, Gail. 1999. "Skepticism, Relevant Alternatives and Deductive Closure." In *Skepticism: A Contemporary Reader*, edited by Warfield & DeRose. Oxford: Oxford University Press.
- Stroud, Barry. 1984. *The Significance of Philosophical Scepticism*. Oxford: Clarendon.
- Vogel, Jonathan. 1990. "Are There Counterexamples to the Closure Principle?" In *Doubting*, edited by M. D. Roth Ross and G, 13–27. Dordrecht: Kluwer Academic Publishers,.
- Wright, Crispin. 1985. "Facts and Certainty." *Proceedings of the British Academy* 71: 429–72.
- . 2002. "(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell." *Philosophy and Phenomenological Research* 65 (2): 330–48.
- Yablo, Stephen. 2014. *Aboutness*. Princeton ; Oxford: Princeton University Press.