

Precis of Assurance

Krista Lawlor

Stanford University

Recent epistemology has focused a lot of attention on the semantics of knowledge ascriptions. An impressive array of theories has been developed (contextualism, invariantism, subject sensitive invariantism, relativism), as well as works that compare and test the alternatives. Simultaneously, an array of analytical theories about what it takes to know (relevant alternatives theory, safety, double safety, practical interest theory, and so on) has been built, often with an eye to fitting a favored semantics of knowledge claims. With all this activity, it is remarkable that almost no attention is given to the purpose of making knowledge claims. What is it one tries to do in claiming to know something?¹

Immediately, the well-trained philosopher may become uneasy: “What is the purpose of asking about the purpose?” It seems we can roll up our sleeves and build semantic and analytic theories without troubling about the function, or normative dimensions of the speech act of claiming knowledge.

One aim of *Assurance* is to demonstrate the value of asking about the function of knowledge claims. The animating idea is that if we start with the normative dimensions of our assurance-giving practice, and give a rational reconstruction of

¹ Semantic and analytic theories include (Cohen (1988), DeRose (2009), Lewis (1996), Rysiew (2001), Stanley (2005), Fantl and McGrath (2009), MacFarlane (2008) For comparative work, see Hawthorne (2004), Blome-Tillman (2009). Craig (1999) is an outlier: he aims to explicate the *concept* of knowledge, by considering its purpose and possible genealogy. My concern is not the concept of knowledge per se.

that practice, we will make semantically and analytically important discoveries about knowledge. This animating idea is in the spirit of J.L. Austin. And Austin is my guide in the ensuing inquiry.

Austin pretty clearly thinks that in general we would be better off in our philosophical endeavors if we kept linguistic usage steadily in view. This much is widely appreciated. Less appreciated is the fact that Austin also has insights into the central preoccupations of contemporary epistemology—about what it takes to know, the meaning of “knows”, and how utterances in natural language come to be true. A second aim of the book is to show how these insights can be forged together to yield a systematic account of knowledge and knowledge claims.

Finally, I aim to show how an Austinian view speaks directly to current problems in epistemology: for instance, by addressing skeptical arguments based on closure principles and on the information-dependence of warrant, while avoiding the shortcomings of standard contextualism. The view also places new problems in the foreground. We haven’t been puzzled enough about how knowledge claims work, or about how they could work, despite the permanent possibility of disagreement and idiosyncrasy in one’s interlocutors. Although skeptical problems occupy a good part of the book, it also aims to deal with the question of how we can manage our epistemic lives together despite our pervasive differences.

The framework

Austin repeatedly encourages us to compare what we do in making first person knowledge claims with what we do in making promises. To have a name for it, I call the act of making a first-person knowledge claim, *assuring*. Austin says that promising and assuring involve a new plunge. In promising to do something, one goes beyond merely saying one will do it. Likewise, in assuring someone of the truth of a claim, one goes beyond asserting it. Promising and assuring involve giving others one’s word, and offering a surety.

Austin stresses that in cases of assuring and promising, the surety one offers is somehow exception-less or “unlimited”, in contrast to the surety one offers with more limited expressions:

... *know* and *promise* are in a certain sense ‘unlimited’ expressions, while... when I guarantee I guarantee that, upon some adverse and more or less to be expected circumstance arising, I will take some more or less definite action to nullify it.²

With a limited expression such as “guarantee”, the speaker may represent a possible failure, and imply she’ll provide a remedy in that event to “nullify” it (“We guarantee against failures of workmanship”). Promises and assurances aren’t like this. A promisor does not represent future failures (“I promise I’ll pay you back, unless I can’t afford to”), even though we all know that promisors fail sometimes. As with promising, in assuring, one does not represent possible failures, even though we all know that assurances sometimes turn out to be unsound.

The comparison with promising ultimately yields a contrast between assurance and assertion. Assertion is a powerful speech act in its own right: with an assertion one vouches for the truth of what one asserts, and one must have reasons to back what one says. In assuring that *P*, one goes further than with an assertion; one must have more than reason to think *P* true, one must have conclusive reasons to think *P* true. The function of assurance is to allow one’s addressee to set aside *all* concerns about defeating possibilities, to provide its recipient with an exclusionary reason to believe *P*.³

² Austin (1946, 173 note).

³ The question of the function of the speech act of assurance might seem to be the obvious focus of the epistemology of testimony, but while theorists have recently offered interesting and varied accounts of how speakers shoulder epistemic responsibility for hearers, it is with attention to the speech act of *assertion*, without recognition of a distinctive speech act of *assurance*. This is true even of “assurance theories” of testimony, tracing back to Moran (2005).

Two questions arise: How do we reconcile the unlimited aspect of “knows” as an expression with fallibilism about knowledge? (By fallibilism, I mean roughly that one’s evidence need only eliminate some ways of going wrong.) And how can one ever think it appropriate to give unqualified surety as one does with “I know” and “I promise”? If claiming to know involves representing oneself as immune from error why on earth would any sane and non-dogmatic person do it? On the receiving end, why would we accept the word of anyone who represented herself as taking on such an apparently extravagant commitment? I call this “the commitment puzzle.”

In order to address these questions, we look for a rational reconstruction of our assurance giving practice. Pressing further with Austin helps us to produce this reconstruction.

One element of the reconstruction concerns the truth conditions for knowledge claims. While Austin does not provide an account of what it is for a knowledge claim to be true, he does make some key remarks about what it is for a statement to be true, where a statement is what is expressed by an utterance. His central idea is that a statement can only be true or false with respect to the situation about which it is made. This might seem anodyne, but in fact the idea has great significance.⁴ Giving an assurance involves claiming knowledge, and knowledge claims are only true in light of the situation the claim is about.

John Barwise and John Etchemendy, whose work on the Liar Paradox makes use of Austin’s idea, give us the following example as illustration:⁵

Claire’s Hand: You are watching a card game, and you note to another observer, ‘Claire has the ace of hearts.’ It so happens that the player before you is not Claire but Abby (you mistake them). Nonetheless at that very instant, Claire herself, playing at a different table, holds the ace of hearts.

⁴ Austin’s insight has been taken up by a number of philosophers and linguists: Barwise and Etchemendy (1987; Barwise and Perry (1998), Travis (2008), Recanati (2004).

⁵ Barwise and Etchemendy (1987, 121).

Your utterance is true or false in light of not just the “descriptive content” of your utterance, namely, *that Claire has the ace of hearts*, but also the situation you are talking about. A situation comprises particular individuals and facts. In Claire’s Hand, the situation includes this particular player, with this particular hand, before you now. Your utterance is true if there is a match between the descriptive content expressed and the situation you are talking about. (Or we may say, your utterance is true if its descriptive content is *true of the situation*.)⁶ Although there is a situation in the world that matches the descriptive content of your utterance, it is not the right situation to make your utterance true, because you’re not talking about it.

Turn now to knowledge ascriptions, and the following illustration.

Bird watching: Alice spies a large bird above the lake. What sort of creature is this? “That’s an osprey”, Ben assures her. “I know it is—look at its white markings.”

To determine whether Ben’s knowledge claim is true, we see whether its descriptive content is true of the situation, S, he is talking about. What is the descriptive content here? “Knows” is an “unlimited” expression, so we can suppose it means something like, “has conclusive reasons to believe the true claim that ...”, or “can rule out all the alternatives to the true claim that...” So the descriptive content Ben expresses with “I know it’s an osprey” is *that is an osprey and I can rule out all the alternatives to its being an osprey*. Ben’s utterance is true if this content is true of the situation, S, that he talks about. The situation, S, Ben is talking about includes both the particular bird he is looking at, and a set of relevant alternatives to the proposition *that it is an osprey*. Let’s stipulate about the case that the situation S includes the relevant alternatives *that the bird is a heron, a hawk or a kingfisher*. (What makes an

⁶ It should be noted that Austinian semantics is quite independent of Austinian speech act analysis. There is no suggestion to the effect that speech acts, or their features, determine the truth conditions of utterances. Rather, an utterance has its “descriptive content” based on the meanings of the individual terms in the utterance, composed in standard compositional ways. And so, for instance, two different speech acts may involve utterances with the same descriptive content.

alternative relevant? Hold that thought—we'll get there.) Suppose also that Ben's evidence is some distinctive white markings, and this evidence eliminates the possibilities that the bird is a heron, a hawk or a kingfisher. Now we see that in S (i) the bird is an osprey, and (ii) Ben's evidence eliminates all the alternatives to its being an osprey. (Ben's evidence eliminates the set {heron, hawk, kingfisher} and that just is eliminating *all* the alternatives in S.) Thus Ben's statement, "I know it is an osprey" is true.

What makes it the case that the situation one talks about is one in which these particular alternatives are relevant? Austin's reminder that we should think about assuring and promising together is important here. When one promises that one will produce some outcome one *represents* oneself as intolerant of failure. That said, it is pretty clear that one's *commitment* in promising does not extend to ruling out all conceivable alternatives in which the outcome does not obtain. The occurrence of some alternatives is irrelevant in judging whether one has lived up to one's commitments. One relieves the promisee of *all* worry by taking care of only *some* eventualities. So the practice of promising has its own problem of relevance. How does it get resolved? The answer suggested by our promising practice is that the relevant alternatives are the alternatives that it is reasonable to hold one responsible for. The parallel between assuring and promising then, suggests a criterion for relevance: reasonableness is what makes an alternative relevant; unreasonableness is what makes an alternative irrelevant.⁷

Here we arrive at the second element of the rational reconstruction of our assurance giving practice. The idea is that an alternative to a target proposition *P* is relevant if a person considering the situation and meeting a *reasonable person standard* would want it eliminated before taking herself to know *P*. What is a reasonable person standard? We can look to legal practice and theory for some guidance here. The law

⁷ Though Austin does not explicitly invoke reasonableness, he does remark upon the need to keep one's challenges to knowledge claims "within reason." He also talks about the need to avoid "silly" or "outrageous" challenges. Austin (1946, 154).

makes extensive use of the concept.⁸ And while there are disputes in legal theory about the content of the standard, there is broad agreement too, for instance: a reasonable person is physically and mentally normal, and has normal world-knowledge, and impartially weighs evidence, for instance, about what precautions to take.⁹

Let me stop for a moment, just to be clear here about argumentative burdens. My claim is not the normative claim that we *should* employ a reasonable person standard in determining epistemic relevance. My claim is that we can give a good rationalizing reconstruction of our practice of making and evaluating assurances, if we think of assurances as being given and received in light of a reasonable person standard. Much of what we do in giving assurance is done as if we employed it.

Conducting our lives together, it matters very much to us to be able to give and receive assurances. Coordination of many kinds is made possible by the attribution of knowledge to oneself and others. But it is a good question how one ever thinks it appropriate to give assurances to another person. There are hurdles one faces. For one thing, one uses an “unlimited” expression “I know” in doing so. Claiming to know involves representing oneself as immune to epistemic failure. We make some progress over this hurdle by marking a distinction between *representing* oneself as immune to failure, and *committing oneself* to eliminating all the relevant alternatives in the situation. The Austinian semantic story allows us to mark this distinction. As in the promising case, one relieves one’s addressee of *all* worry by taking care of only *some* eventualities. The presence of a reasonable person standard—a standard to which one has some access in virtue of common world-knowledge and a less than super-human capacity for impartiality—permits one to use one’s own take on what

⁸ “As everyone is well aware, ... [there is] a very general tendency in the law to rely upon the standard of reasonableness as a criterion of right decision-making, of right action, and of fair interpersonal relationships within the law of property, the law of obligations, and family law. Even as a few illustrative examples, they suffice to ground the thesis that reasoning about reasonableness is a matter of great moment within the operations of the law.” MacCormick (1999, p.3)

⁹ Gilles (2001).

is reasonable to worry about in order to gauge whether one is in any position to take up the commitments of assurance giving. A second hurdle in assurance giving is epistemic: one is sometimes in the dark about the idiosyncratic needs and practical interests of one's interlocutors. Appeal to a shared standard of reasonableness permits us to negotiate our ignorance about each other. My completing my birder's "Life List" may depend on a correct identification of the bird on the lake, but you don't need to know that to give an assurance based on your grasp of what is generally reasonable to worry about in the way of alternatives. Clearly there is more to be said about how this works. For instance a closer look at our practice reveals that there is a default assumption that a speaker can use her own sense of what is reasonable to worry about, unless a hearer explicitly defeats that assumption with an explicit statement of what rests on the claim. If my life depends on identifying the bird above the lake (!), I need to speak up.

The combination of Austinian semantics for knowledge attribution, and an Austin-inspired account of relevance in terms of reasonableness, jointly rationalizes our assurance giving practice. A sane and non-dogmatic person can use an "unlimited" expression such as "knows" in part because the truth of her knowledge claim depends on what is *relevant in the situation talked about*, and in part because what counts as relevant in the situation will be *something she has access to*, something she can use her own judgment about.

Applications

The Austinian view has some advantageous features. The contrast between Austinian semantics and standard semantic contextualism helps to highlight a few of these features.

Standard semantic contextualism holds that the meaning or semantic contribution of "knows" to the proposition expressed varies with the context of utterance.¹⁰ What

¹⁰ DeRose (2009), Cohen (1988). In an alternative formulation, standard semantic contextualism bypasses the claim that the *meaning* of "knows" depends on context, and holds instead that *the proposition* expressed by a knowledge claim depends on

affects the context is the epistemic standards of the attributor. Standard semantic contextualism nicely explains how we find apparent variance in truth-values of knowledge claims despite sameness of the subject and her evidence. But contextualism has problems too.

(i) *Indirect Reports*. An osprey lights on the lake. Ben the birdwatcher rightly claims to know that the bird is an osprey. Otto the ornithologist also rightly claims to know the bird is an osprey. Not much is at stake for Ben, but a lot rides on a correct identification for Otto. Otto's standards are more demanding and "know" has a concomitantly stronger meaning in his mouth than in Ben's. In such a case, according to standard contextualism the meaning of "knows" in Ben's mouth is something weaker and more forgiving than in Otto's mouth.¹¹ The propositions each express with "I know it's an osprey" differ. Ben expresses, roughly, *that it is an osprey, and my evidence meets ordinary standards*; Otto expresses, roughly, *that it is an osprey, and my evidence meets high standards*. Otto's report, "Ben said he knows it's an osprey", involves a more demanding meaning than Ben himself uses. How can Otto report Ben's knowledge claim, then? This is a problem for standard contextualism. Austinian semantics addresses the problem: the meaning of "knows" and the propositional content (the "descriptive content") of knowledge claims is stable across these utterances, despite different situations of utterance. The descriptive content in Ben's utterance and Otto's utterance is *that the bird is an osprey and I can eliminate all the alternatives to its being an osprey*. So Otto can report Ben's utterance "Ben says he knows it's an osprey" without misrepresenting the content of Ben's utterance.

(ii) As can be seen from the case of indirect reports, Austinian semantics shares some features of the view John MacFarlane calls "non-indexical contextualism."

context. Not all contextualisms are "standard" by my lights: for instance, Neta (2003) articulates contextualism about evidence, and avoids the problems for standard contextualism that I highlight below.

¹¹ As noted, alternatively, standard contextualism avoids context-sensitivity of meaning but adopts context-sensitivity of proposition expressed. Either way, contextualism runs into difficulty with indirect reports.

Consequently, it provides a response to the problem of “faultless disagreement” that is like the non-indexical contextualist response.¹² Disputants need not be talking past each other when disagreeing about knowledge or matters of taste, since their dispute focuses on the same claim.

(iii) *Meaning Deflation/Vacuous Knowledge*. Standard contextualism runs into two problems in addressing closure-based skepticism. Closure-based skepticism, familiarly, arises this way. If a closure principle of knowledge under known implication holds, then if one knows by looking across the lake, say, *that the bird is an osprey*, one also knows what follows (if one recognizes the implication), *that it is not a decoy*, left on the lake by a careless hunter. But it seems extraordinary that one could know such a thing on the basis of looking across the lake; and so neither does one know that the bird is an osprey. Standard contextualism, familiarly, addresses the problem by restricting closure to a fixed context. Epistemic standards affect the context, and the meaning of “knows.” So, in a relaxed standards context R, our subject Ben knows-by-relaxed-standards *that the bird is an osprey*, as well as *that it is not a decoy*, and in the more demanding context H, Ben fails to know-by-high-standards *that it is not a decoy*, and so also fails to know-by-high-standards *that the bird is an osprey*. There is no contradiction here, so no paradox, and no threat to ordinary knowledge.¹³

There are two related problems with this standard contextualist approach to the paradox. First, the problem of meaning devaluation¹⁴: Ben knows *that it is not a decoy* in the ordinary setting, not because he has evidence to this effect, but because “knows” means something suitably weak. How weak? Weak enough so that Ben can

¹² MacFarlane (2008). I say more about faultless disagreement in “Austinian semantics and non-indexical contextualism” (manuscript).

¹³ For ease, I state the paradox in material mode, and talk of *knowing propositions*, not of the *truth and falsity of “knowledge” claims*. Doing so does not affect the issues under discussion here—that is, unless the contextualist wants to insist that some knowledge claims (“I know it’s not a decoy”) *cannot* be assessed in the relaxed context; but then one needs a story about why one cannot assess the claim with respect to this context.

¹⁴ Kornblith (2000).

stand in the knowledge relation to the proposition *that it is not a decoy* while lacking evidence specifically about whether there have been hunters about using decoys, and so forth. Second, the problem of extraordinary knowledge¹⁵: If we allow context to remain fixed while considering new propositions, knowledge is too easy to come by, as it does not require an appropriate contribution from the subject's evidence.¹⁶ The birdwatcher, Ben, says, "I know it is an osprey" truly in R, and he deduces and believes *Q, that it is not a decoy left by a hunter*. Then when Ben says, "I know it is not a decoy left by a hunter" his claim is also true in the context R.¹⁷

The Austinian account diagnoses our intuitive sense that Ben comes by this purported knowledge too easily: Ben's evidence (the distinctive markings on the bird's throat, say) is enough to eliminate the relevant alternatives to *P*, but not enough to eliminate the relevant alternatives to *Q*. And Austinian semantics avoids both the foregoing problems. First, it avoids semantic devaluation: "knows" means nothing less than "has conclusive reasons to believe the true claim that ..." and it means this on every occasion of its use. Second, Austinian semantics blocks extraordinary knowledge: there is no holding fixed a context and allowing the evaluation as true or false any propositions whatsoever with respect to that fixed context. This is a key departure from standard contextualism. The standard semantic contextualist envisions a single unchanging context in which Ben knows both that it's an osprey and not a decoy. Austinian semantics does not implement unchanging contexts fixed by epistemic standards. As new propositions are considered, new alternatives are relevant; the situation talked about is changed.

I have briefly touched on some advantages of Austinian semantics, as against standard semantic contextualism. The Austinian framework more broadly helps to address a range of problems and puzzles through appeal to the notion of reasonableness. For instance, appeal to reasonableness allows for a principled

¹⁵ Heller (1999, 206–7) calls it "vacuous knowledge."

¹⁶ Some argue that appearances notwithstanding, the birdwatcher can know *Q* by deducing it from *P*; see Stine (1999), Klein (1995). I find such responses implausible.

¹⁷ This is so, even if the effect of his utterance is also to shift the context to the more demanding *H*. His claim is nonetheless true as assessed in *R*. (See note 13).

response to Crispin Wright's skeptical argument from information dependence, and the articulation of a position that admits Liberalism about perceptual justification without Dogmatism.¹⁸

The Austinian view, with its account of relevance as reasonableness, provides the resources to defend a principled fallibilism about knowledge.¹⁹

¹⁸ The terms *liberal* and *dogmatist* are from Pryor (2004).

¹⁹ Many thanks to Laura Schroeter, David Hills and Ram Neta for helpful comments.

Works cited

- Austin, J. L. 1946. "Other Minds." *Supplement to the Proceedings of the Aristotelian Society* 20: 148–87.
- Barwise, Jon, and John Etchemendy. 1987. *The Liar: An Essay on Truth and Circularity*. Oxford University Press, USA.
- Barwise, Jon, and John Perry. 1998. *Situations and Attitudes*. Center for the Study of Language and Inf.
- Cohen, Stewart. 1988. "How to Be a Fallibilist." In *Philosophical Perspectives*, edited by Tomberlin, 2:581–605.
- Craig, Edward. 1999. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press, USA.
- DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism and Context*. Vol. 1. Oxford: Oxford University Press.
- Fantl, Jeremy, and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press, USA.
- Gilles, Stephen. 2001. "On Determining Negligence: Hand Formula Balancing, the Reasonable Person Standard, and the Jury." *Vanderbilt Law Review* 54 (3): 813–61.
- Heller, Mark. 1999. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77 (2): 196–208.
- Klein, Peter. 1995. "Skepticism and Closure: Why the Evil Genius Argument Fails." *Philosophical Topics*, Epistemology, 23 (1): 213–36.
- Kornblith, Hilary. 2000. "The Contextualist Evasion of Epistemology." *Philosophical Issues* 10 (1): 24–32. doi:10.1111/j.1758-2237.2000.tb00004.x.
- Lewis, David. 1996. "Elusive Knowledge." *Australian Journal of Philosophy* 74: 549–67.
- MacCormick, Neil. 1999. "Propter Honoris Respectum: Reasonableness and Objectivity." *Notre Dame Law Review* 74 (1575).
- MacFarlane, John. 2008. "The Assessment Sensitivity of Knowledge Attributions." In *Epistemology: An Anthology*, edited by Kim Sosa. Oxford: Blackwell.
- Moran, Richard. 2005. "Getting Told and Being Believed." *Philosophers' Imprint* 5 (5).
- Neta, Ram. 2003. "Contextualism and the Problem of the External World." *Philosophy and Phenomenological Research* 66 (1): 1–31.
- Pryor, James. 2004. "What's Wrong with Moore's Argument?" *Philosophical Issues* 14 (Epistemology): 349–78.
- Recanati. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35 (4): 477–514.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. 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.

Travis, Charles. 2008. *Occasion-Sensitivity: Selected Essays*. Oxford University Press, USA.