

The Express Knowledge Account of Assertion

Forthcoming in the *Australasian Journal of Philosophy**

John Turri
Huron University College
john.turri@gmail.com

Abstract: Many philosophers favor the simple knowledge account of assertion, which says you may assert something only if you know it. The simple account is true but importantly incomplete. I defend a more informative thesis, namely, that you may assert something only if your assertion expresses knowledge. I call this ‘the express knowledge account of assertion’, which I argue better handles a wider range of cases while at the same time explaining the simple knowledge account’s appeal. Section 1 introduces some new data that a knowledge account of assertion well explains. Section 2 explains the simple knowledge account’s advantage over two of its main competitors. Section 3 presents a problem for the simple account and offers a solution, which is to adopt the express knowledge account. Section 4 encapsulates the case for the express knowledge account, and offers a unifying vision for the epistemology of belief and assertion. Section 5 answers an objection. Section 6 briefly

* This is a preprint of an article whose final and definitive form will be published in the *Australasian Journal of Philosophy*, which is available online at: <http://www.tandf.co.uk/journals/>.

sums up. Even those who ultimately reject my conclusion can still benefit from the new data presented in section 1, and learn an important lesson from the problem discussed in section 3, which demonstrates a general constraint on an acceptable account of the norm of assertion.

1. The Simple Knowledge Account of Assertion

Increasingly popular these days is the **simple knowledge account of assertion** ('the simple knowledge account' for short), which says you may assert Q only if you know Q.¹ The 'may' here designates epistemic permission—knowledge is essentially the requisite epistemic relation for assertion. It does not follow, nor is it suggested, that you morally or prudentially may assert Q only if you

¹ I call it the 'simple knowledge account' rather than the 'knowledge account' because it is the simpler of the two knowledge accounts discussed here. Notice I *don't* use 'simple' the same way Timothy Williamson [2000: 241–2] does. In any ordinary sense, both knowledge accounts discussed here are simple. But in Williamson's sense, whereas the simple knowledge account does count as simple, the express knowledge account arguably doesn't.

Williamson [2000: Chapter 11] presents the most influential case for the simple knowledge account. See also [Unger 1975: Chapter VI; DeRose 2002; Reynolds 2002; Hawthorne 2004: Chapter 1.3]. I omit evidence involving lottery cases, primarily because I don't think it's very good evidence, because I don't think it's out of line to assert that you'll lose the lottery (compare [Williamson 2000: 246]). I also omit discussion of the nature of the rule identified by the knowledge account, which is often said to be *the* 'constitutive rule' of assertion uniquely. Some commentators, e.g. [Kvanvig 2009], express no interest in whether knowledge is the constitutive norm of assertion, but only in whether knowledge is a norm of assertion.

know Q.

The simple knowledge account's proponents adduce a wide range of linguistic data that it well explains. This section reviews five previously noticed data points, and introduces two new ones in the process. I acknowledge that competing views purport to explain some of this data equally well [e.g. Weiner 2005, Douven 2006, Lackey 2007, Hill and Schechter 2007, and Kvanvig 2009]. But I won't pause to evaluate these claims, because my primary purpose in this paper lies elsewhere, namely, to motivate a new and improved version of the knowledge account.

First, when I assert Q, even if Q has *nothing* to do with me or what I know, asking me 'How do you know?' is normally appropriate [Unger 1975: 263–4, acknowledging Michael Slote for the insight]. The simple knowledge account nicely explains this. By asserting Q, I represent myself as having authority to do so, rendering your question appropriate.

Second, asking 'How do you know?' is understood as implicitly challenging my authority to make the assertion. More aggressive than asking 'How do you know?' is 'Do you really know that?' [Williamson 2000: 252–3]. More aggressive yet is 'You don't know that!' or 'You don't know what you're talking about'. The simple knowledge account explains this range of aggressiveness. 'How do you know?' implicitly challenges my authority to assert Q, by asking me to demonstrate that I do in fact have it; 'Do you really know that?' explicitly challenges my authority, by questioning whether I have it; and 'You don't know that!' explicitly rejects my authority.

Third, when someone asserts Q to us, if it turns out that he doesn't know Q, we're entitled to resent his assertion. Moreover if he reasonably believes or, worse, knows that he doesn't know Q, then we feel that his assertion was dishonest [Unger 1975: 260–3].² The simple account nicely explains both of these things. If he doesn't know Q but asserts it anyway, then he misrepresents himself as having the authority to do so, which misrepresentation we may resent. And if he *knows* that he doesn't know Q but asserts it anyway, then he *knowingly* misrepresents himself, which is dishonest.

Fourth, when someone asks you whether Q, even when Q has nothing to do with you or your knowledge, normally you may appropriately respond by saying 'I don't know'. The simple knowledge account makes it easy to understand why. You're informing this person that you lack the authority to directly answer her question. We expect people to answer only if they know [Reynolds 2002: 140].

Fifth, instead of responding to a prompt by saying 'I don't know', you could without loss have instead responded with 'I can't tell' or 'I can't say'. The simple knowledge account explains the interchangeability of these locutions. 'Tell' and 'say' both mean 'assert' (as is typical in ordinary speech), and 'can' expresses the concept of permission (again, typical). And since knowledge is the norm of assertion, to lack permission *just is* to lack knowledge.

Sixth, assertions of the form 'Q but I don't know that Q' strike

² Unger considers the latter case to be "of a piece with lying."

us as inconsistent [Unger 1975: 256–60; Williamson 2000: 253–4]. Their *content* is perfectly consistent, so whence the inconsistency?³ The simple knowledge account can explain why. If the simple knowledge account is true, then you may make that assertion only if you know each conjunct: *Q* on the one hand, and *I don't know Q* on the other. So your knowing the first conjunct would falsify the second conjunct, in which case you couldn't possibly know the conjunction. But in asserting the conjunction, you represent yourself as knowing it (because you represent yourself as having authority to do so), which explains the inconsistency. What you assert is inconsistent with how you represent yourself.

Seventh, it's all too common to hear things like 'I can tell you that your paper is still out with the referee'.⁴ Consider how odd it would sound to say 'Your paper isn't still with the referee, but I can tell you that it is' or 'I don't know whether your paper is still with the referee, but I can tell you that it is'. (We get the same effect if we replace 'I can tell you' in these utterances with '(what) I can say (is)'). The simple knowledge account nicely explains the oddity of those assertions. The second conjunct states that I have authority to assert that your paper is still out with the referee. But the first conjunct either directly denies that I have the authority, in the case of 'I don't know', or obviously entails that I lack it, in the case of 'your paper isn't still with the referee'.

³ As an anonymous referee pointed out, such assertions are consistent only if *Q* itself is consistent.

⁴ Thanks to E.J. Coffman for conversation on this point.

2. Anything You Can Do

The simple knowledge account's competitors include the the **truth account**, which says you may assert Q only if Q is true [Weiner 2005], and the **reasonable belief account**, which says you may assert Q only if you reasonably believe Q [Hill and Schechter 2007, and Kvanvig 2009].⁵

To the extent that we approach consensus on anything in epistemology, reasonable true belief is necessary but not sufficient for knowledge.⁶ Because reasonable true belief is necessary for knowledge, the simple knowledge account forbids anything forbidden by either the truth account or the reasonable belief account. But since reasonable true belief doesn't suffice for knowledge, the simple knowledge account forbids assertions that neither the truth account nor the reasonable belief account forbid.

Given these logical points, it will come as no surprise that proponents of the simple knowledge account argue that these two competitors fail to forbid obviously impermissible assertions.⁷ On the

⁵ Douven [2006] and Lackey [2007] advocate something similar to the reasonable belief account, but they don't require belief for proper assertion; they require only that it be reasonable or rationally credible for you to believe what you assert.

⁶ Sartwell [1992] and Hetherington [2002] dissent.

⁷ E.g. [Williamson 2000: Chapter 11, esp. 242, 260–2]. The simple knowledge account's critics will argue that it forbids obviously permissible assertions, including Gettiered assertions [e.g. Weiner 2005: 230–1; Hill and Schechter 2007: 109; Lackey 2007: 596, 598 ff; Levin 2008: 369–70; Kvanvig 2009]. I acknowledge the point, but I'll set it aside here because my goal is not to defend the simple knowledge account against this objection (see [Williamson 2000: 256] and [DeRose 2002:

one hand, many people judge that false assertions are automatically improper. But the reasonable belief account doesn't forbid these, because some false beliefs are nevertheless reasonable. On the other hand, many people likewise judge that an assertion expressing a random guess (a *random assertion*) is automatically improper, regardless of whether the assertion luckily turns out to be true. But the truth account doesn't forbid lucky random assertions.

By contrast, the simple knowledge account forbids both false assertions and lucky random assertions. Proponents of the simple knowledge account thus claim a significant advantage over their opponents.⁸

3. I Can Do Better

Does the reasonable belief account forbid a lucky random assertion? Initially you might think that the answer is obviously 'yes', because a lucky guess doesn't qualify as a reasonable belief. But the answer is actually 'no'. Working through this problem for the reasonable belief account will reveal a more general problem facing extant accounts of assertion, including the simple knowledge account, which is my ultimate focus here.

A case will demonstrate the point. Suppose Randy reasonably

180, 199 n. 23] for what have become standard responses.)

⁸ I recognize that proponents of these competing views have responses to the criticisms rehearsed here. My purpose in this paper allows me to simply bypass these dialectical complexities without unfairness.

but non-occurrently believes Q. Neither his belief nor his evidence for it is presently conscious. Oddly, this remains so even as his partner turns and asks whether Q. (Despite its oddity, this is surely possible.) Unable presently to recall, Randy randomly guesses that Q is true, and confidently expresses this random guess to his partner. The reasonable belief account doesn't forbid *this* random assertion. But it's clearly defective.

In response to such cases, proponents of the reasonable belief account might modify their theory, so that it becomes the **modified reasonable belief account**, which says: you may assert Q only if (i) you reasonably believe Q and (ii) you assert Q at least partly because it is reasonable for you to believe Q.⁹ But this modification isn't enough to forbid lucky random assertions. To see why, imagine that Randy's partner asked him whether Q because Randy told her yesterday that Q, but she can't now clearly remember whether he said 'Q' or 'not-Q'. Moreover—and this is the key—Randy told her Q yesterday *because* he had good evidence to believe Q (the same good evidence he's retains till now, but doesn't presently recall). None of this changes the fact that Randy's assertion is clearly defective. But the modified reasonable belief account doesn't forbid it.

The best response to the problem is to require that the assertion *express* a reasonable belief, that is, the best solution is to revise the account to say: you may assert Q only if your assertion *expresses* a

⁹ Compare [Lackey 2007: 608]. Again, I note that Lackey denies that you must believe Q to properly assert Q.

reasonable belief that Q. Call this the **express reasonable belief account**. Randy asserts something that he reasonably believes, but his assertion fails to express that reasonable belief, which explains its impropriety.

Proponents of the modified reasonable belief account might complain that although Randy asserts Q because it is reasonable for him, the fact that it's reasonable does not explain his assertion *in the right way*, so the case does not motivate the move to the express reasonable belief account.¹⁰ But then the question simply becomes 'Well, what's *the right way*?'. I say the right way is for the assertion to express the reasonable belief. So we're right back to the express reasonable belief account.

Now we can appreciate that the simple knowledge account faces an analogous problem, namely, that we can easily imagine possible cases where someone knows Q, and asserts Q, but nonetheless the assertion is defective. Two examples follow.¹¹

Suppose Randy remembers that Q, and thereby knows that Q, but presently this knowledge is non-occurrent. Neither his knowledge nor his evidence for it is presently conscious. Oddly, this re-

¹⁰ Matt Weiner helpfully suggested this as a possible response, without positively endorsing it.

¹¹ Lackey [2007: 609–610] considers similar examples and comes close to appreciating their full significance, but proposes something like condition (ii) of the modified reasonable belief account to handle them. Levin [2008: 374–5] and Brown [forthcoming] also both argue that the simple knowledge account fails to forbid impermissible assertions. But as Coffman [unpublished] demonstrates, their arguments mistake prudential or moral impermissibility for epistemic impermissibility. My discussion poses a completely independent challenge to the simple knowledge account.

mains so even as his partner turns and asks whether Q. (Despite its oddity, this is surely possible.) Without actually trying to recall, Randy randomly guesses that Q is true, and confidently expresses this random guess to his partner. The simple knowledge account doesn't forbid *this* random assertion, because Randy knows what it is that he's asserting. But the assertion is clearly defective.

Spiro is a spiteful guy who relishes causing people emotional pain. Out of spite, he plans to tell Lois that her fiancée just died. Some time before he embarks to execute his plan, he receives a text message from a reliable informant reporting that Lois's fiancée has indeed just died. So Spiro knows that the fiancée died. But this knowledge doesn't motivate him in the least to tell Lois that her fiancée died. He goes ahead and tells her *out of pure spite*. Intuitively this assertion is epistemically defective. (It is also morally defective, but that's a different point.) And yet he asserts something that he knows, so the simple knowledge account doesn't forbid it.

The best response to such cases resembles the earlier shift from the reasonable belief account to the expression of reasonable belief account. That is, the best response is to require that the assertion *express knowledge*. The **express knowledge account of assertion** says you may assert Q only if your assertion *expresses* knowledge that Q. Randy asserts something he knows, but his assertion fails to express his knowledge, which explains why it's defective. The same goes for Spiro. Cases like these give a considerable advantage to the express knowledge account over the simple knowledge account.

4. Beyond the Simple Knowledge Account

Expressing knowledge of Q requires knowing Q. Thus the express knowledge account forbids everything that the simple knowledge account forbids. And the express knowledge account can equally well explain the linguistic data from section 1. But the express knowledge account can also explain why some assertions are defective *despite* the fact that the speaker knows the proposition asserted. It's not enough that you know Q when you assert Q. It's not even enough that you assert Q because you know Q. The assertion must also express knowledge.¹²

The express knowledge account also has at least one other very attractive feature. It fits seamlessly into a powerful and increasingly popular approach in epistemology, namely, virtue epistemology [e.g. Zagzebski 1996, Greco 2007, Sosa 2007]. On this view, knowledge is true belief manifesting cognitive virtue. When coupled with the express knowledge account, this produces an elegant and unified picture of the primary individual and social dimensions of knowledge, namely, its acquisition and communication. Knowledge

¹² If asked for an analysis of this concept, I would provisionally say that your assertion expresses mental state M just in case M non-deviantly causes your assertion. *Expression* here is a special way for concrete token mental states to *manifest* themselves. So in the relevant sense, to express a belief or knowledge requires that you be in the relevant token mental state. There's a sense in which you can 'express the belief' that Q even if you don't believe Q. But this just means that you declaratively utter 'Q'. (Here 'belief' just means *the propositional content*, as when we say 'You and I share the same belief'.) Thanks to an anonymous referee and Aidan McGlynn for conversation on these points.

is true belief manifesting virtue, and proper assertion is speech manifesting knowledge.

For those who think that knowledge isn't a mental state and so can't really be expressed in the way my account requires, there's an easy fix. Understand 'expresses knowledge that Q' to mean 'expresses a belief that Q, in virtue of which you count as knowing that Q', or 'expresses a knowledgeable belief that Q'.

The concept of *expression* is familiar to us from everyday thought and speech. Invoking it doesn't beg any questions in the present context. It already features prominently in philosophical discussions of ethics, action and belief. And the express knowledge account is interesting and substantive even absent an analysis of *expression*. So it's fair for us to appeal to it here, without defending a theory of it.

5. Having Authority to Assert *versus* Permissibly Asserting

This section answers an objection to my view, imagined to come from a defender of the simple knowledge account. My answer helps to further motivate the express knowledge account. The objection which I'll answer emerges against the backdrop of a standard response to a popular criticism of the simple knowledge account.

The popular criticism of the simple knowledge account is that it's too strong. It's too strong because it's acceptable to assert things that you reasonably, but falsely, believe to be true, whereas you can't know something false. In response, Timothy Williamson

[2000: 257] suggests that such assertions are *reasonable* but nevertheless *impermissible*. You might reasonably but falsely believe Q, in which case your assertion would be reasonable despite being impermissible, just as you might reasonably but falsely believe that you were obeying the speed limit, in which case your driving speed would be reasonable despite being impermissible. On this way of looking at it, the popular criticism mistakes a reasonable assertion for a permissible one. Keith DeRose [2002: 180] makes a related distinction between *primary* and *secondary* propriety in following a rule. Primary propriety consists in your actually following the rule. Secondary propriety consists in its being reasonable for you to think that you're following the rule. A secondarily but not primarily proper assertion is blameless despite being impermissible. So on DeRose's way of putting it, the popular criticism mistakes a blameless assertion for a permissible one.

Now the objection to my view goes like this.¹³ For simplicity I focus on Randy's case alone, the application to Spiro's being exactly similar. Perhaps proponents of the simple knowledge account might pursue a related strategy to handle the intuition that Randy impermissibly asserts. They might argue that Randy's assertion is *permissible despite being unreasonable* (or blameworthy), and I have simply mistaken an unreasonable (or blameworthy) assertion for an impermissible one.

The proponent of the simple knowledge account thus marshals the same distinction in response to my objection, as he marshalled

¹³ Thanks to an anonymous referee for suggesting the objection.

just a moment ago in response to the other objection. But whereas a moment ago he accused one critic of mistaking a reasonable assertion for a permissible one, he's now accusing me of mistaking an unreasonable assertion for an impermissible one.

I have two points in response. First, this manoeuvre places even more weight on what many consider to be the weakest link in the overall defence of the simple knowledge account—namely, leaning on the permissible/unreasonable distinction, and accusing opponents of bungling it when considering particular cases—and so comes at a cost. If there were no alternative, it might be worth paying this cost, in light of the simple account's explanatory benefits. But there is an alternative, the express knowledge account, which not only handles Randy's case effortlessly, but also preserves the spirit of the simple knowledge account. So why wouldn't we want to trade in the simple account for my alternative? What do we gain from resisting the move to the express knowledge account?

The second part of my response to this objection is to emphasize that I don't deny that Randy *has the authority* (or permission) to assert what he does. (Again I focus on Randy's case for simplicity.) To know Q is to have authority to assert Q, so Randy obviously has the authority. What I deny is that he *permissibly asserts* Q. Having authority to assert isn't the whole story of permissibly asserting. To the extent that someone intuits that Randy's assertion is permissible, it might well be because they recognize that Randy has the authority to assert what he does, which could be true *even if Randy impermissibly asserts*. This fits into a more general pattern,

as I'll now show.

In general, permissibly A-ing (where 'A' names an act type) requires more than A-ing while you have authority to A. Despite having authority to A, you might nevertheless impermissibly A. For instance, suppose that in virtue of having evidence E, it's permissible for you to believe that Q. Despite this, you're not permitted to believe Q in just any old way. In such a case, permissibly believing Q requires you to properly base your belief on evidence E. Believing based on entirely unrelated grounds won't suffice—that would be a bad belief, irresponsibly formed [Turri 2010]. Likewise an executioner might be authorized to kill the prisoner. But that doesn't mean it would be permissible for the executioner to fatally beat the prisoner with a tire iron before he has his last rites. That would be an impermissible killing. Permissibly killing the prisoner requires the executioner to kill the prisoner in the appointed way. In general, when we may A, it's because we have permission to A in some specific way.¹⁴ Unsurprisingly, then, permissibly asserting requires more than asserting something which you have permission to assert. Indeed this is precisely what we should have expected.

We might put the matter this way. The difference between the simple knowledge account and the express knowledge account is the difference between stating the requirement of *having the authority to assert*, and stating the requirement of *permissibly asserting*.

¹⁴ There is a limit to the specificity, of course, but that doesn't spoil the point, and the details needn't concern us here.

It's perhaps worth noting that my point here closely relates to another view congenial to those sympathetic to the simple knowledge account of assertion. 'Belief,' Williamson says, 'aspires to knowledge' [2000: 1]. Indeed, 'the point of belief is knowledge.' Williamson also claims that 'believing p stands to asserting p as the inner stands to the outer,' so that the simple knowledge account of assertion 'corresponds to the norm that one should believe p only if one knows p' [2000: 255–6]. But notice, to say that knowledge is *the point* of belief, or that belief *aspires* to knowledge, goes far beyond saying that you should believe something only if you know it. Rather, it's more like enjoining you to believe only if your believing *counts* as knowing. A belief that doesn't count as knowledge isn't good, even if it's belief in something that you do in fact know. And just as the point of belief is knowledge, not merely to believe something you know, it stands to reason that the point of assertion is to express knowledge, not merely to say something that you happen to know. It seems, then, that the express knowledge account of assertion coheres better with the overall thrust and tenor of the Williamsonian vision of the relationship between assertion, belief and knowledge.

6. Conclusion

In closing, let me reiterate what I hope to have accomplished here, and what I have expressly not attempted to accomplish, beginning with the latter. I have not defended the simple knowledge account

from prior objections. I have not addressed attempts to undermine the cumulative explanatory case for the simple knowledge account. By contrast, I have introduced some new data that a knowledge account well explains, and I have argued that the express knowledge account outperforms the simple knowledge account. I also hope to have highlighted an important general challenge facing multiple existing accounts of assertion, not just the simple knowledge account. The discussion in section 3 teaches us a general lesson that transcends the evaluation of the simple knowledge account, namely, that a fully satisfactory account of the norm of assertion will include something like *an expression condition*.¹⁵

¹⁵ For helpful conversation and feedback, I thank Matt Benton, E.J. Coffman, Geoffrey Cork, Aidan McGlynn, Ernest Sosa, Angelo Turri, Matt Weiner, and two anonymous referees.

REFERENCES

- Brown, Jessica forthcoming. Knowledge and Assertion, *Philosophy and Phenomenological Research*.
- Coffman, E.J. unpublished. Two Mistakes about Epistemic Propriety.
- DeRose, Keith 2002. Assertion, Knowledge and Context, *Philosophical Review* 111/2: 167–203.
- Douven, Igor 2006. Assertion, Knowledge, and Rational Credibility, *Philosophical Review* 115/4: 449–85.
- Greco, John 2007. The Nature of Ability and the Purpose of Knowledge, *Philosophical Issues* 17: 57–69.
- Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hetherington, Stephen 2002. *Good Knowledge, Bad Knowledge*, Oxford: Oxford University Press.
- Hill, Christopher S. and Joshua Schechter 2007. Hawthorne's Lottery Puzzle and the Nature of Belief, *Philosophical Issues* 17: 102–122.
- Kvanvig, Jonathan L. 2009. Assertion, Knowledge, and Lotteries, in *Williamson on Knowledge*, ed. Duncan Pritchard and Patrick Greenough: 140–160.
- Lackey, Jennifer 2007. Norms of Assertion, *Noûs* 41/4: 594–626.
- Levin, Janet. 2008. Assertion, Practical Reason, and Pragmatic Theories of Knowledge, *Philosophy and Phenomenological Research* 76/2: 359–84.
- Reynolds, Steven L. 2002. Testimony, Knowledge, and Epistemic Goals, *Philosophical Studies* 110/2: 139–161.
- Sartwell, Crispin 1992. Why Knowledge Is Merely True Belief, *Journal of Philosophy* 89/4: 167–180.
- Sosa, Ernest 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, v. 1, Oxford: Oxford University Press.
- Turri, John 2010. On the Relationship Between Propositional and Doxastic Justification, forthcoming in *Philosophy and Phe-*

nomenological Research 80/2.

Unger, Peter 1975. *Ignorance: A Case for Skepticism*, Oxford: Oxford University Press.

Weiner, Matthew 2005. Must We Know What We Say? *Philosophical Review* 114/2: 227–251.

Williamson, Timothy 2000. *Knowledge and Its Limits*, Oxford: Oxford University Press.

Zagzebski, Linda 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.