

Revisiting Norms of Assertion

John Turri

Philosophy Department and Cognitive Science Program
University of Waterloo
200 University Avenue West
Waterloo, Ontario N2L3G1, Canada
john.turri@gmail.com

Abstract: A principal conclusion supported by convergent evidence from cognitive science, life science, and philosophy is that knowledge is a central norm of assertion — that is, according to the rules of the practice, assertions should express knowledge. That view has recently been challenged with new experiments. This paper identifies a critical confound in the experiments. In the process, a new study is reported that provides additional support for the view that knowledge is a central norm of assertion.

Keywords: assertion; knowledge; communication; norms

Words: 2095

Introduction

A challenge facing any communication system is that the interests of sender and receiver often diverge, leading to dishonest signaling, such as false predator alarm calls. If dishonesty proliferates too much, then the signals will eventually be ignored and the communication channel rendered worthless. Stable and enduring communication systems thus include mechanisms that promote honest signaling. In humans, assertion is a principal means of communicating information. What prevents humans from dishonestly asserting enough to destabilize the practice?

I have argued that the human practice of assertion is at least partially sustained by a socially policed rule that *assertions should express knowledge*. This view attempts to shed light on human communication by placing it in a broader context of scientific understanding of animal communication. In particular, on this view, the human practice of assertion is partly sustained by mechanisms similar to those that sustain non-human communication systems. This provides the view with a deep and principled theoretical motivation. Importantly, it is also supported by convergent evidence from several areas of cognitive science, including developmental and cross-cultural psychology (Turri & Park in press; for a comprehensive review, see Turri 2016a; for a shorter review, see Turri 2017a).

Markus Kneer (2018) recently challenged this view. He assessed four hypotheses from the philosophical literature on assertability, which differ on whether to substitute an appropriate

phrase pertaining to belief, justified belief, truth, or knowledge into the schema, “Assert that P only if ___,” where “P” stands for a proposition. Kneer focused his critical energy on two of these views: “Assert that P only if you know that P” and “Assert that P only if P is true.” He conducted several studies in which participants read a brief scenario and judged whether a specific proposition was true, assertable, and known. Kneer manipulated (between-subjects) whether the relevant proposition was false or true in the scenario. Replicating previous findings (see below for references), he found that a nontrivial number of participants in several key conditions judged that an agent should assert a proposition that is false and not known. He interprets this as strong evidence against a normative connection between knowledge and assertion. By contrast, he interprets some of his results as suggesting that “justified belief” is the norm of assertion (“Assert that P only if you have a justified belief that P”).

Theory

Before proceeding to the main critical point, I would first like to emphasize some theoretical points of agreement. Despite being framed principally as a critical response to my work, Kneer’s contribution is arguably best understood as a response to a superficially similar but fundamentally different view from the theoretical philosophy literature. That view is only loosely related to empirical evidence. It assumes, among other things, that there is a single, exceptionless standard of assertability (for discussion, see Turri 2014; see also footnote 3 and the General Discussion of Kneer 2018). On this approach, if assertions should express knowledge, then in no circumstance

should anyone assert anything that they do not know. This view has been repeatedly rejected in the literature (Turri 2014; Turri 2016a: p. 63; Turri 2017a), because it is inconsistent with how social rules tend to work and also with a range of experimental findings. Social rules tend to tolerate exceptions. You should donate to charity, but you do not violate this rule by not donating today, this week, or this month; parents should pay attention to their children, but parents typically do not violate this rule by sleeping. Similarly, while existing evidence shows that knowledge is a central norm of assertion, it does not support the conclusion that knowledge is an exceptionless standard of assertability, or that all people tend to reliably link assertability with knowledge in all circumstances. Instead, existing evidence shows that there is a detectable, often very strong, central tendency to link judgments about assertability to knowledge. Better understanding this central tendency's limitations is a principal objective of ongoing research, to which Kneer's paper contributes.

Just as "should" does not express an exceptionless standard, neither does it identify the unique standard of assertability. People distinguish assertions that should not be made from blameworthy assertions (Turri 2013; Turri & Blouw 2015), so in that sense there are at least two standards. This reflects a more general fact that people reliably distinguish what someone *should not do* from what he is *blameless for doing*. For instance, people reliably distinguish between broken promises that should have been kept and those that were blamelessly broken (Buckwalter & Turri 2015; Turri 2016c; Chituc, Henne, Sinnott-Armstrong & De Brigard 2016; Turri 2017b). People also distinguish beliefs that should not be held from blameworthy beliefs (Turri 2015). Similarly, Kneer's findings provide some initial evidence that people distinguish assertions that

should be made from assertions that can be *justified*.

Confound

People are more likely to attribute knowledge when the response options contrast knowledge to ignorance, such as “knows/doesn’t know,” than when they contrast knowledge to a different mental state, such as “really knows/only believes,” “really did know/thought she knew,” or “actually knows/only thinks he knows” (e.g. Cullen 2010; Buckwalter 2014). A likely explanation is that the *plain* options (“knows/doesn’t know”) increase knowledge attribution because some participants answer based on how things seem to the agent. In comparison, the options involving a *contrast* with how things seem to the agent (e.g. “actually knows/only thinks”) do not force participants to choose between reporting *how things actually are* and *how things seem to the agent*. Similar observations have been made about probing for assertability attributions with the options “should/shouldn’t” compared to “actually should/only thinks he should” (Turri 2016b).

Unfortunately, Kneer mixed the two types of options (plain/contrast) across the two probes (knowledge/assertability). For example,

Should Bob say “Jill drives an American car”?

- Yes
- No

Does Bob really know that Jill drives an American car?

- Yes, Bob really knows that Jill drives an American car.
- No, Bob merely believes that Jill drives an American car.

Thus when Kneer reports observing significantly higher rates of agreement with the first ques-

tion (over 60% selected “Yes”) compared to the second (only about 20% selected “Yes ...”), there are multiple explanations for the disparity. One explanation, favored by Kneer, is that most people think that the agent actually should assert a proposition he does not actually know. Another explanation is that the results are confounded by unbalanced response options: the knowledge question was asked in a way that tends to reduce the rate of attribution, whereas the assertability question was not.

To help evaluate these two explanations, I conducted a simple experiment to directly test the effect of response options, while also varying the statement’s truth-value. Two hundred and one U.S. residents were recruited on Amazon Mechanical Turk and randomly assigned in a 2 (truth-value: false, true) × 2 (option type: plain, contrast) between-subjects design. The truth-value factor varied which version of this scenario participants read:

(False/true) Bob has a friend, Jill, who has driven a Buick for many years. A Buick is an American car. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill [replaced it with a Mercedes, which is a German car / regularly cleans its chrome rims with a non-abrasive cloth]. ¶¹ One day Bob’s wife asks him, “Does your friend Jill drive an American car?”

The false version of the story is taken verbatim from Kneer’s study, and the true version is a close control matched for length and complexity. The option factor varied the options offered for the assertability and knowledge items:

¹ Indicates a paragraph break on the participant’s screen

Bob _____ say that Jill drives an American car. (assertability)

Bob _____ that Jill drives an American car. (knowledge)

The plain options were “should not/should” and “doesn’t know/knows.” The contrast options were “only thinks he should/actually should” and “only thinks he knows/actually knows.” Participants always rated assertability on the first screen, then went to a new screen and rated knowledge. Finally, all participants went to a new screen and answered a comprehension question: “Is it true that Jill drives an American car?” (response options “yes/no”). All response options were randomly rotated. Ninety-two percent of participants (185 of 201) passed the comprehension question.

Binary logistic regression revealed a very large effect of option for each attribution (assertability, knowledge), qualified by a significant interaction between option and truth-value (false, true). (See Tables 1 and 2.) Visualization of the results shows that the interaction is due to the difference between options having the predicted effect in the false condition, dramatically reducing attribution of both assertability and knowledge and, indeed, reversing the central tendency of both from agreement to disagreement. (See Figure 1.) By contrast, the difference between options had a much smaller effect in the true condition, with agreement remaining the strong central tendency for both attributions regardless of option type.

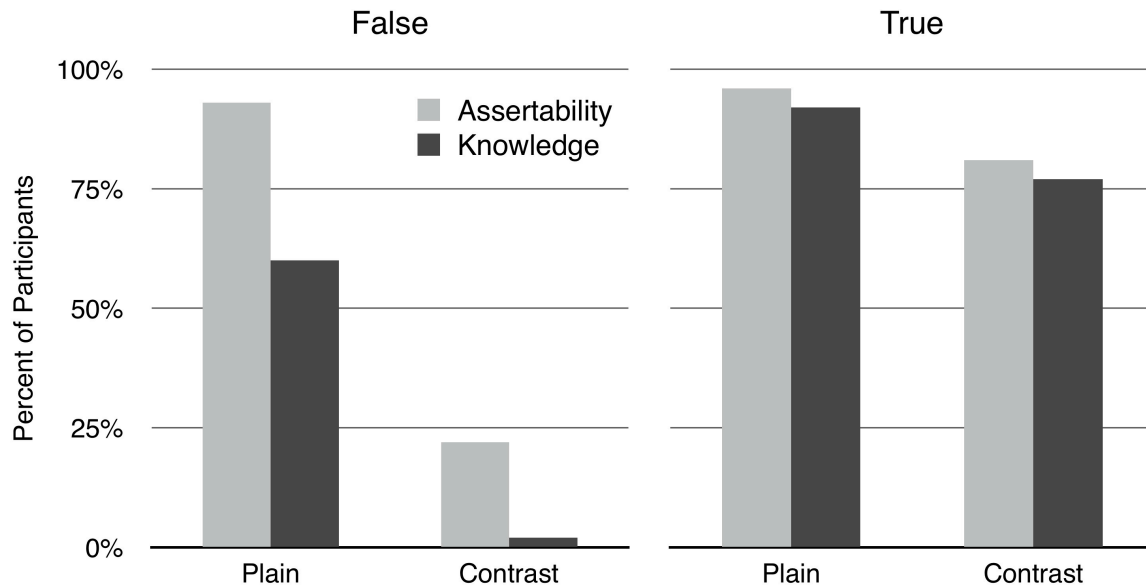


Figure 1. Percent of participants attributing assertability and knowledge (within-subjects), broken down by whether the relevant proposition was false or true (between-subjects) and using either plain or contrast options (between-subjects).

Table 1. Logistic regression predicting assertability attributions.

	B	SE	Wald	df	p	Odds Ratio	Odds Ratio 95% CI	
							LLCI	ULCI
Truth-value	0.61	0.94	0.43	1	.513	1.85	0.29	11.61
Option	-3.82	0.70	29.89	1	<.001	45.50	11.58	178.80
Truth-value × option	2.11	1.07	3.87	1	.049	0.12	0.02	0.99
Constant	2.57	0.60	18.33	1	<.001	13.00		

Note: $\chi^2(3, n = 185) = 81.48, p < .001$, Cox & Snell $R^2 = .356$, Nagelkerke $R^2 = .520$. Reference class for truth-value: false. Reference class for option: plain.

Table 2. Logistic regression predicting knowledge attributions.

	B	SE	Wald	df	p	Odds Ratio	Odds Ratio 95% CI	
							LLCI	ULCI
Truth-value	5.00	1.07	21.89	1	<.001	148.00	18.25	1200.50
Option	4.17	1.06	15.50	1	<.001	64.71	8.12	515.70
Truth-value × option	2.94	1.23	5.72	1	.017	0.05	0.01	0.59
Constant	-3.78	1.01	14.00	1	<.001	0.02		

Note: $\chi^2(3, n = 185) = 104.71, p < .001$, Cox & Snell $R^2 = .432$, Nagelkerke $R^2 = .583$. Reference class for truth-value: false. Reference class for option: plain.

These results replicate previous findings supporting a normative connection between knowledge and assertion (e.g. Turri 2016b), and they further demonstrate the importance of using consistent response options across the two types of judgment, especially when, contrary to how things seem to the agent, the relevant proposition is false. For example, focusing on the false conditions in this sample, we find that 93% of participants attribute assertability using the plain options, but only 2% attribute knowledge using the contrast options (a much larger disparity than even Kneer observed). Ignoring the difference between response options, we might conclude that the vast majority of participants dissociate assertability and knowledge in this case. But if we had instead probed for knowledge using plain options, we would have found that 60% also attributed knowledge. Or if we had instead probed for assertability using contrasting options, we would have found that only 22% attributed assertability. Either way, provided that we use a consistent probing method, very different conclusions are warranted. Overall in the false conditions in this study, 28% of participants attributed assertability without attributing knowledge (collapsing across the option factor, which did not affect the rate at which participants did this, Fisher’s exact

text, $p = .149$). This is comparable to the rate reported in previous studies (23% – 30%) that have been interpreted, including by Kneer, as supporting a normative connection between knowledge and assertability (Turri 2013; Turri 2016b).²

Indeed, Kneer did not report a single study in which he probed for knowledge and assertability in the same way (see the Appendix of Kneer 2018).

Conclusion

In summary, first, Kneer and I agree that we should not accept the assumptions that there is a unique norm of assertion, or that the norm imposes an exceptionless necessary condition. Second, as the results of a new experiment show, Kneer's studies were critically confounded by a difference in response options across different probes. Correcting for that serious error, the results replicate previous findings that support, even by Kneer's lights, a normative connection between assertion and knowledge.

In conclusion, I commend Kneer for providing interesting initial evidence that what should be asserted differs from what one is "permitted" to assert, according to the rules of the practice. The relevant between-subjects comparison in the false-belief condition of his first experiment appears tightly matched and responsibly reported. Future work could profitably investigate how

² For example: "It appears that roughly 3/4 of people respond as a factive account would predict, whereas roughly 1/4 of people respond consistently with the non-factivity intuition" (Turri 2013: 285). "The rate at which participants linked knowledge and assertability remained constant across both sorts of procedure (~75%)" (Turri 2016b: 285).

permissibility relates to blamelessness, which prior work has also shown to differ from what should be asserted. Similarly, future work should investigate the function that the different norms play in stabilizing and sustaining communication practices. This issue has been discussed theoretically (Turri 2017d), but it remains to formulate credible empirical tests of the relevant theoretical predictions. Finally, future work should not re-import false, counterproductive assumptions from the philosophical literature.

Acknowledgments — For helpful feedback and discussion, I thank [[omitted at review stage]]. This research was supported by the Canada Research Chairs program and the Social Sciences and Humanities Research Council of Canada.

References

- Buckwalter, W. (2014). Factive verbs and protagonist projection. *Episteme*, 11(4), 391–409. <http://doi.org/10.1017/epi.2014.22>
- Buckwalter, W., & Turri, J. (2015). Inability and Obligation in Moral Judgment. *PLOS One*, 10(8), e0136589. <http://doi.org/10.1371/journal.pone.0136589.g004>
- Chituc, V., Henne, P., Sinnott-Armstrong, W., & De Brigard, F. (2016). Blame, not ability, impacts moral “ought” judgments for impossible actions. *Cognition*, 150(C), 20–25. <http://doi.org/10.1016/j.cognition.2016.01.013>
- Cullen, S. (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1(2), 275–296. <http://doi.org/10.1007/s13164-009-0016-1>

- Kneer, M. (2018). The norm of assertion: empirical data. *Cognition*.
- Turri, J. (2013). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, 129(2), 279–291. <http://doi.org/10.1016/j.cognition.2013.06.012>
- Turri, J. (2014). Knowledge and suberogatory assertion. *Philosophical Studies*, 167(3), 557–567. <http://doi.org/10.1007/s11098-013-0112-z>
- Turri, J. (2015). The radicalism of truth-insensitive epistemology: truth’s profound effect on the evaluation of belief. *Philosophy and Phenomenological Research*, 93(2), 348–367. <http://doi.org/10.1111/j.1933-1592.2002.tb00206.x>
- Turri, J. (2016a). *Knowledge and the norm of assertion: an essay in philosophical science*. Cambridge: Open Book Publishers. Retrieved from <http://www.openbookpublishers.com/product/397/knowledge-and-the-norm-of-assertion--an-essay-in-philosophical-science>
- Turri, J. (2016b). Knowledge and assertion in “Gettier” cases. *Philosophical Psychology*, 29(5), 759–775. <http://doi.org/10.1080/09515089.2016.1154140>
- Turri, J. (2016c). Compatibilism and incompatibilism in social cognition. *Cognitive Science*. <http://doi.org/10.1111/cogs.12372>
- Turri, J. (2017a). Experimental work on the norms of assertion. *Philosophy Compass*, 12(7), e12425. <http://doi.org/10.1002/ajp.22097>
- Turri, J. (2017b). How ‘ought’ exceeds but implies ‘can’: Description and encouragement in moral judgment. *Cognition*, 168, 267–275. <http://doi.org/10.1016/j.cognition.2017.07.008>
- Turri, J. (2017d). Sustaining rules: a model and application. In J. A. Carter, E. C. Gordon, & B. Jarvis (Eds.), *Knowledge first*. Oxford University Press.

Turri, J., & Blouw, P. (2015). Excuse validation: a study in rule-breaking. *Philosophical Studies*, 172(3), 615–634. <http://doi.org/10.1007/s11098-014-0322-z>

Turri, J., & Park, Y. (in press). Knowledge and assertion in Korean. *Cognitive Science*.