

Iffy Predictions and Proper Expectations

Synthese 191 (2014): 1857–1866

Please cite the published version, available at link.springer.com: [here](#)

Matthew A. Benton

University of Oxford

John Turri

University of Waterloo

Abstract: What individuates the speech act of prediction? The standard view is that prediction is individuated by the fact that it is the unique speech act that requires future-directed content. We argue against this view and two successor views. We then lay out several other potential strategies for individuating prediction, including the sort of view we favor. We suggest that prediction is individuated normatively and has a special connection to the epistemic standards of expectation. In the process, we advocate some constraints that we think a good theory of prediction should respect.

1 Introduction

The category of *assertives* is a broad and important category of speech acts. The category's namesake, the speech act of assertion, has received the lion's share of attention in recent years. Debates about assertion's nature and norms have loomed large in both the philosophy of language and epistemology (e.g. Geach 1965, Williamson 2000: ch. 11, Turri 2011, Benton 2011, the essays in Brown & Cappelen 2011, and Turri 2013b). But assertion isn't the only interesting and important assertive that has received attention in the literature recently. For example, recent work focuses on the speech acts of predicting and guaranteeing (Benton 2012, Kelp 2013, Turri 2013a). And earlier work on assertion itself has often relied on important assumptions about assertion's relationship to other assertives, again fo-

cusing on prediction and guaranteeing (e.g. Weiner 2005 and Turri 2010, respectively).

The category of assertives includes guessing, conjecturing, hedged asserting, flat-out asserting, swearing and guaranteeing, among others. According to the standard view (Searle 1979, pp. 12–13), at least two important features bind these together into a useful genus. First, their illocutionary point is to commit the speaker, to some degree, to the truth of some claim.¹ Second, they have a word-to-world direction of fit, in that their purpose is to represent the way the world is.² But what individuates the species of assertives within the broader genus? In particular, what might individuate *prediction* specifically, such that a speech act is a prediction rather than a guess or a flat-out assertion?

One might think that the very idea that there is a speech act of prediction distinct from (say) flat-out assertion requires argument. Here are some reasons to think there is a distinct speech act of prediction, even though one can predict by asserting: first, in some conversational contexts, it is clear to the participants that an outright declarative utterance ‘P’ is not a flat-out assertion, but rather a prediction. For example, with the election race very tight, when one is asked, ‘Whom do you predict will win the next election?’, one’s outright utterance ‘Obama will win’ plau-

¹More precisely, the point is to commit the speaker *either* to the truth of some claim, or to the untruth of some claim. *Denial* is also an assertive, the twin of assertion, but it commits the speaker to the untruth of the claim denied. We can set this complication aside.

²It might be thought that, as a corollary of the two distinguishing features just mentioned, by performing an assertive speech act, one thereby expresses some degree of belief, or at least implies that one believes to some degree that the relevant proposition is true (compare what Searle says at 1979, pp. 12–13 in light of 1979, p. 4). But this doesn’t seem essential, since *guessing that P* does not, even by default, express or imply that the guesser believes, even to some degree, that P is true. Even Searle seems to recognize this when he remarks, ‘The degree of belief and commitment may approach *or even reach zero*’ (1979, p. 13, emphasis added). This is reflected in the fact that the following isn’t Moore-paradoxical: ‘I’m not the least bit confident that P is true, but still, my guess is that P’.

sibly merely predicts rather than (outright) asserts; likewise, were one prompted in such a context with ‘Whom would you say will win?’, ‘Whom do you expect will win?’, or (perhaps most commonly) ‘Whom do you think will win?’, by answering ‘Obama’ one will naturally be regarded as merely predicting rather than outright asserting that Obama will win. One of the tell-tale signs of this is that in such a case, by uttering ‘P’, the speaker does not represent herself as knowing P, and perhaps in some contexts does not even represent herself as outright believing P.³ Second, in such a case it will often seem strange for the conversational participants to respond to such utterances with questions like ‘How do you know?’, since it is clear to those involved that the speaker is merely predicting rather than asserting.

Thus before beginning in earnest, we’d like to flag a simplifying device. It’s not uncommon to perform one speech act by performing another, as when one makes a request by asking a question (as with ‘Can you pass the salt?’), or swears to the truth of something by asserting it (as with sworn testimony). And arguably one can perform speech acts non-verbally too, at least once the speech act has been established in the community. We don’t deny that one can predict by guessing or even flat-out asserting; nor do we deny that one can predict non-verbally. But we do think it’s best to focus our investigation here on cases of direct, verbal prediction—that is, on cases where one predicts that P verbally, but doesn’t do so by guessing that P or asserting that P (e.g. by the schema ‘I predict that P’). This simplifies the task of discerning what, if anything, constitutes the speech act of prediction. Nuances involving indirect or nonverbal prediction can then be added as the need or interest arises.

³Such occasions will be likely be ones in which the speaker somehow signals (perhaps through intonation or hesitancy, etc.) that she is less confident. But in cases without such signals, an outright utterance (even if a mere prediction) will tend to represent the speaker as believing, as is suggested by the fact that while ‘How do you know?’ will seem out of place, ‘Why do you think that?’ or ‘Why do you believe that?’ will be acceptable questions or challenges. (Thanks to an anonymous referee here.)

2 The iffy view

Let's begin with what is, by an overwhelming margin, the most common answer we encounter to our question. When we ask, 'What individuates prediction within the class of assertives?', the first response tends to be, 'Why, *futurity*, of course!' More specifically, the view is that the speech act of prediction is individuated by the fact that it alone must have a future-directed content: a speech-act is a prediction if and only if it must have future-direct content.⁴ Because this view says that the speech act of prediction is individuated by future-directed content, we'll call it *the iffy view* of predictions. Searle appears to advocate the iffy view when he writes,

The differences, for example, between a report and a prediction involve the fact that a prediction must be about the future whereas a report can be about the past or present. (1979, p. 6)

The iffy view gains intuitive support from etymology. Our word 'predict' derives from the Latin *praedicere*, which is composed of the prefix *prae-*, which means 'before', and *dicere*, which means 'say'. Thus 'predict' arguably means 'say before' (or 'foretell'), and 'say before' is most plausibly interpreted as 'say before it happens'. It's a benefit of the iffy view that it can uphold this *etymological desideratum*, as we might call it, in this way.⁵

A minor adjustment to the iffy view is required to handle predictions about the future made in the simple present tense. The historical present

⁴We say 'future-directed' rather than 'future-tensed,' because, first, predictions can be made without using the future tense (e.g. 'I am going to be at your party', which uses the present progressive). And second, according to many linguists, many languages (including English) do not even have a future tense; yet it still seems plausible to say that predictions can be made in such languages. We'll also assume that the iffy view holds that the future-directed contents must also be *contingent*.

⁵Another iffy view builds the future-directedness into the speaker's intention to communicate something about (what the speaker takes to be) the future; the iffy views canvassed later in the main text also could be construed at the level of speaker intention. Such a view would handle what seems right about a subject who wrongly thinks it is 2010 and predicts that P about 2011. But such *intentionally iffy* approaches fall prey to some of the obvious counterexamples to be discussed below.

is typically characterized as the use of the present tense when describing past events. For example, when explaining the demise of the dinosaurs to a group of children, we could say, 'It's sixty-five million years ago. It seems like just another day to the enormous reptiles roaming the land. But little do they know that disaster is about to strike, in the form of a giant meteor.' A similar device, what we might call the *antecedent present*, allows us to use the simple present tense to describe future events. For example, in the early twenty-first century, a climatologist might describe the likely effects of global warming in the late twenty-first century by saying, 'The polar caps completely melt, the ocean level rises, and coastal cities worldwide are devastated.' To take another example, if asked to predict the outcome of next week's playoff series between Boston and Vancouver, you could say, 'Boston wins the series in five games,' and this would count as a prediction that Boston will win the series in five games. To accommodate the antecedent present, the iffy view should be adjusted to read: predictions are individuated by requiring future-directed content, even if this content isn't expressed in a distinctively future-tensed way.

But even when understood in the most plausible way, the iffy view faces a serious problem: it excludes many paradigm cases of prediction. Theories and theoreticians predict things, and sometimes they predict things that have happened long ago. Big Bang Theory and its proponents predict that the early universe was extremely hot, that galaxies used to be much closer together, and much else besides about the past, even unto its utmost reaches.⁶ Or to take a more mundane case, suppose we all know that the ball game is over and has a winner. We can still sensibly ask you, 'Whom do you predict won the game?', and you could sensibly answer, 'I predict that Boston won.' Future-directed content is not essential to pre-

⁶Physicists really do talk this way. Google searches reveal the following examples: 'If we trace the rate of expansion and extrapolate in time, we predict that the uni-verse in the past was smaller than the universe of today'; and 'Cosmologists predict that the early Universe was full of small galaxies which led short and violent lives.'

diction. The iffy view is false.

3 The modified iffy view

A natural response to this last objection begins by revisiting the etymological desideratum. Earlier we were happy to understand ‘predict’ as meaning ‘say before’. But say before *what*? We assumed that it meant *say before it (the predicted event) happens*. But as we just saw, this is unacceptable because it rules out paradigm cases of predictions about past events. Another interpretation of ‘say before’ is that predictions pertain to *the predictor’s future evidential situation*. On this view, ‘I predict that Boston won’ is elliptical for something like ‘I predict that *we’ll discover that Boston won*’, and ‘I predict that the early universe was extremely hot’ is elliptical for something like ‘I predict that *we’ll find evidence best explained by the hypothesis that the early universe was extremely hot*’. In short, ‘say before’ amounts to ‘say before you learn’.⁷ Call this the *modified iffy view*. According to the modified iffy view, when properly understood, the predictions we adduced as counterexamples to the iffy view do have future-directed contents: they are about what the predictor will learn in the future about the truth of the proposition in question. Understood this way, they are not counterexamples to the modified iffy view.

The primary problem with the modified iffy view is that it mischaracterizes what our predictions are about. When we make the prediction, ‘Boston will win the series in five games,’ it doesn’t seem to us that we’re making a prediction about what we’ll learn about the outcome of the series. Rather, it seems to us that the truth of our prediction depends solely on the outcome of the series, not what we learn about the series. Consider also this example. Late in life, Edwin Hubble lobbied for the Nobel

⁷We intend ‘learn’ to be understood flexibly. Variations on this proposal could feature ‘know’, ‘discover’, ‘acquire good evidence’, etc. These details won’t figure into our evaluation of the view, so we note them here briefly, only to set them aside.

Prize Committee to make work in astronomy eligible for the Nobel Prize in physics. At the time of Hubble's death, the Committee hadn't re-classified work in astronomy. Suppose that on his deathbed Hubble said, 'I predict that the Committee will make work in astronomy eligible for the Nobel Prize in physics.' Hubble's prediction isn't falsified by the fact that he died before learning that the Committee did re-classify astronomy. Rather, it seems that his prediction is made true by the fact that the Committee re-classified astronomy, regardless of whether Hubble lived to learn that fact.

4 The hybrid iffy view

It might be thought that a hybrid of the iffy and modified iffy views can handle all the cases. On this *hybrid iffy view*, predictions always have future-directed contents, but not necessarily the same type of future-directed content. When the event in question hasn't yet happened (or been completed), the prediction is about the future event directly. This handles cases like Hubble's deathbed prediction. But when the event has already happened, the prediction is about the predictor's future epistemic state. This handles cases like the Big Bang theorist's predictions about the early universe, and the prediction about who won the game last night. We would be happy if the hybrid iffy view enjoyed the benefits of each of the earlier iffy views, but the drawbacks of neither.

But the hybrid iffy view runs afoul of intuitions about the following kind of case.

(Forgetful Fans) You and I both know that Vancouver is scheduled to play Boston this week, but we've both forgotten whether the game was to be played today (in two hours' time, say), or whether it in fact was played yesterday at that time. (Neither of us has come into contact with any sports news in the last few days.) Without settling the matter of what day

the game was actually scheduled for, you ask, ‘Any predictions about the game?’ I respond, ‘My prediction: Vancouver wins.’

It seems clear that in *Forgetful Fans*, regardless of whether the game was yesterday or is still to be played today, my prediction will be about the same thing. It is a serious strike against a view if it can’t respect this *similarity desideratum*, as we might call it. And the hybrid iffy view can’t respect the similarity desideratum. If the game is in fact today in two hours, the hybrid view rules that the prediction is genuinely about the future game; by contrast, if the game was played yesterday, it rules that the pre-diction is about what we will find out about who won the game. We find this result unacceptable.⁸

5 Predictable expectations

Timothy Williamson (2000) has argued influentially that assertion is individuated normatively, to wit, assertion is the unique speech act such that knowledge is its unique constitutive norm. Williamson’s thesis has proved useful as a basis for articulating norms of other assertive speech acts (e.g. Kelp 2013, Turri 2013a). In light of this, one might try to individuate prediction similarly, by identifying one or more norms unique to it. Given that mere predictions have less assertoric force than flat-out assertions, it stands to reason that prediction’s constitutive norm would be something weaker than knowledge. Here we sketch an approach in this spirit.

⁸A different way of preserving the core of the iffy view is to combine it with a kind of *speech act externalism* (speech acts are ‘wide’): factors blankly external to the individual can make an illocutionary difference. According to one way of developing this *externalist iffy* account, you simply fail to make a prediction when you’re blind to the fact that the event you’re speaking of isn’t in the future. Although we’re not opposed to speech act externalism (see Davis 1994 and Gauker 2003; but compare Harnish 2009), its application in the present context seems *ad hoc*, and it doesn’t afford enough respect to the similarity desideratum.

Predictions can plausibly be individuated on the broad spectrum of assertives by appealing to the mental state which predictions seem to express, namely (mere) *expectation*. We don't think this ordinary term needs explication for the present proposal to be taken seriously, but here is a sketch of the notion we have in mind: (mere) expectation is a mental state of slight commitment, which requires regarding a proposition as *more likely than not*, where this is a matter of having a higher credence in that proposition than in its negation. One's mere expectation that P will be 'proper' when one's credences are apportioned to one's evidence, namely when one's evidence makes P more likely than not-P. (Belief and knowledge also entail commitment; but they go beyond *slight* commitment, so they don't involve *mere* expectation.⁹) Roughly, the amount of commitment required for expectation exceeds that of suspicion, but falls short of outright belief. This might not uniquely pick out expectation, but it helps narrow the field considerably. This also suggests that *suspicion* might set the normative standard of *conjecture*, in the way that *expectation* does for *prediction*.

Thus the view we want to put forth is the following:

(EXP) *One may predict that P if and only if one properly expects that P.*

We shall first consider the necessity direction, then the sufficiency direction.

On this view, the norm of predicting is the necessary condition: *Predict that P only if one properly expects that P*. Since fulfilling this condition entails expecting P, a quick test for this norm is a cancelling conjunction:

⁹Perhaps *mere* expectation falls significantly short of *full* expectation, and perhaps full expectation suffices for outright belief; at the very least, this creates room for differences between mere prediction and fully confident prediction. This notion expectation is somewhat akin to talk of credences in formal epistemology, but because our main interest is in how one can express such commitment, we have little to say about that here.

(1) I don't expect that P, but still, I predict that P is/will be true.

(2) I predict that P, but I don't expect that P.

If (EXP) is true, then predicting P in the absence of expecting that P ought to sound infelicitous because improper. Uttering (1) or (2) in the indicative mood does not sound great to us, in that it opens one up to the retort: 'Well why are you predicting that P rather than not-P?' But (1) and (2) don't exactly clash either: one could, it seems, expect neither that P nor that not-P, and nevertheless still use the 'I predict that' construction to predict. Such a prediction might still seem somewhat improper in the absence of expectation; but for those who judge it to be perfectly acceptable, it would be nice to have an explanation of why it can seem so.

That one might utter (1) or (2) acceptably can be explained if (EXP) is coupled with an account that distinguishes predictions from other speech acts by *what they earn you*: one earns credit for predictions in a way that one does not earn credit for guesses or conjectures. Likewise, one earns more credit for successful use of speech-acts as one rises through the spectrum of assertives that are more forceful in their commitment: one earns more credit (and puts more on the line) for successful hedged assertions than for predictions, more for outright assertions than for hedged ones, and more for guarantees than for outright assertions (cf. Turri 2010, pp. 84–86). Because expecting, and predicting, ought to be based on evidence or reasons for thinking that a proposition is (more likely than not to be) true, one's competence gets rewarded with a small amount of credit when one expresses expectation through a true prediction. An account which combines the normative and credit elements has the resources to say why (1) and (2), though somewhat acceptable, don't sound great: (1)'s first conjunct, and (2)'s second conjunct, effectively concede that one doesn't have the proper reasons favoring P. Because it discloses that the speaker doesn't really expect what is predicted, it lessens the credit she could have accrued

for a straight-up prediction; as such, the second conjunct seems effectively *downgraded* to a mere guess.¹⁰ Because guesses sit so closely on the spectrum to predictions, they seem to pattern in similar ways; but the notion of credit earned and extracted, and the corresponding reward for having competently based one's expectation on reasons, helps make sense of the difference between predictions and guesses.¹¹

Now for the *sufficiency* direction of (EXP), which gives the condition for being well-enough positioned for predicting: if one (appropriately) expects that P, then one is in a position to predict that P. Otherwise put: *If one properly expects that P, then one may predict that P.*¹² Given this, predic-

¹⁰We don't say that it no longer counts as a prediction, but rather that its prospects have been eviscerated to the point where one has undermined the motivation for making a prediction rather than a guess. Something similar may be happening in cases where one claims to predict something which one does not regard as more likely than not: e.g., before the NCAA basketball tournament, in which 64+ teams participate, one 'predicts' that a given team will win it all, even though one's proper credence that they will win is not more than 0.5. In these special cases, no one possesses proper expectation of the kind we outline above, and when it is understood by all involved that no one could have that, what passes for proper expectation shifts to the alternative which seems *more likely than any other*. In such contexts, it is unclear whether one is in fact still predicting rather than guessing; indeed, as an anonymous referee points out, it may be partly a matter of cultural convention whether, at such margins, one labels a speech act a mere guess or instead a prediction (albeit one *without* reasonable expectation).

¹¹If one doubts that guesses differ from predictions in this way, compare the following conjunctions:

(3) I have no idea, but I will/am going to go ahead and guess that P.

(4) I have no idea, but I will/am going to go ahead and predict that P.

(4) sounds a bit worse to us than (3), and we suspect it is because one should have some leaning or reason, however slight, on which to base a prediction, whereas guesses don't require even that.

¹²One obvious difficulty here is closure: we expect lots of possibilities each of which are individually more likely than not but which taken together are not more likely than not (e.g. we will expect that the next throw of the 6-sided fair die won't come up 1; but we expect that individually of each number on the die, even though we know that some number must turn up). If (EXP) is correct, won't it sanction a speaker in predicting multiple outcomes that are, taken together, not possible? We acknowledge this difficulty but note that similar closure issues arise for assertion, evidence, justification, belief, and knowledge (though different examples will apply). So (EXP) faces no special problem

tion is the unique speech act such that its minimal sufficient permissibility condition is the minimal sufficient permissibility condition for expectation.

The following observations provide some evidence for the proposed equivalence of minimal sufficient permissibility conditions. We're unwilling to say that the standards of expectation and prediction come apart. It strikes us as absurd to say, 'That was not predictable, though it was to be expected', or, 'That was to be expected, but it wasn't predictable.'¹³ Our unwillingness to allow the standards of expectation and prediction to come apart is well explained by the fact that they *don't* come apart—they're one and the same. Whenever something is predictable, it is also reasonable to expect it; and whenever it is reasonable to expect something, it is also predictable.

(EXP)'s sufficiency condition also coheres well with the spectrum of commitment for assertives: being positioned to engage in an assertive speech act makes one also well-enough positioned (epistemically speaking) to engage in each of the weaker speech-acts on the spectrum. Thus

- (i) If one appropriately *expects* that P, then one is in a position to predict that P and also in a position to guess that P.

And the other assertives can be similarly situated by other sufficiency claims:

- (ii) If one *knows* that P, then one is in a position to assert that P, and to assert that one believes that P, and to predict

here.

¹³These strike us as akin to other absurd claims relating speech acts to their normative standards, such as the following about knowledge and assertion: 'You could have known that it was going to happen, but you were in no position to say whether it would happen', 'She knows whether it's true, but she is in no position to say whether it's true.'

that P, and to guess that P.¹⁴

- (iii) If one *knows that one knows* P, then one is in a position to guarantee that P, and to assert that P, and to assert that one believes that P, and to predict that P, and to guess that P.

Clearly, the sufficiency claims (i)–(iii) go beyond what (EXP) itself is committed to; but they unify the assertive speech-acts on the scale of commitment, while helping to show how they are individuated. Most relevant for our purposes, (i) reveals how close guesses are to predictions (which gives a clue to why they can seem to pattern together). But (i) also permits an asymmetry between them: being in a position to guess doesn't necessarily suffice for being in a position to predict. Otherwise put, *guessing* requires no evidence or credence, and earns no credit; *predicting* requires some evidence and expectation, earns partial credit, and implies being positioned to guess; *asserting* requires knowledgeable belief, earns full credit, and implies being positioned to predict; and *guaranteeing* requires knowing that one knows, earns full credit *plus* bonus marks, and implies being positioned to assert.

One virtue of (EXP), and the sufficiency claims (i)–(iii) above, is that they well explain the overlap between asserting and predicting: on the one hand, one can predict by way of asserting precisely because asserting also commits one to proper expectation; but on the other hand, (EXP) delineates what constitutes the speech-act of mere prediction, and how one could engage in prediction, and properly so, without outright asserting. Furthermore, the (EXP) account allows one to forego building into the speech-act of prediction a stricter norm according to which one should predict that P only if one does *not know* that P. Including such a norm is

¹⁴Though see Brown 2010 and Lackey 2011 for possible reasons to doubt this sufficiency thesis.

tempting because (mere) predictions tend, and are typically expected by hearers, to be made by speakers *without* knowledge of what is predicted. But the sufficiency condition of (EXP) provides the resources for explaining how it is that speakers typically do not know when issuing predictions: given (i)–(iii), not knowing is a pragmatic fact about the conditions under which we typically opt to predict because, when one merely expects that P, prediction is the strongest speech-act to which one can commit given one’s (perceived) epistemic position. And it is plausible to suppose that there is a speech-act equivalent to the ‘Assert the Stronger’ rule,¹⁵ on which: if you’re going to perform an assertive speech-act, you should perform the strongest assertive that your evidence permits (and no stronger). Such a rule would explain, first, why hearers can typically expect that someone who opts to predict (rather than assert, or guarantee) doesn’t know what is predicted, and second, why it would seem objectionable for one to use a weaker speech act such as predicting (or conjecturing, or guessing) when one clearly knows. (Note that such a rule isn’t specific to prediction, so it cannot be used to individuate prediction.)¹⁶

6 Conclusion

The standard view is that the speech act of prediction is individuated by future-directed content. Given the serious problems we have identified for it, the standard view seems unsustainable, in both its original and modified forms. A much more promising approach individuates assertion in terms of either the mental state it expresses or its distinctive evidential

¹⁵On which see DeRose 2009, pp. 86–87.

¹⁶There may be further fruitful ways of distinguishing prediction. It may be plausible to distinguish prediction from other weak assertives such as guesses or conjectures by the way it relates to other aspects of discourse, similar to how ‘I reply that P’ differs from ‘I object that P’. Prediction may also be distinguishable stylistically, similar to how announcing that P differs from confessing that P. We do not take up such highly nuanced approaches here.

norms. We have defended the view that the mental state of *expectation*, and the evidential norms associated with it, is well suited to fill this role. We hope, and expect, that this is a very promising first step toward a fully adequate account of prediction.¹⁷

References

- Benton, Matthew A. 2011. "Two More for the Knowledge Account of Assertion." *Analysis* 71: 684–687.
- . 2012. "Assertion, Knowledge, and Predictions." *Analysis* 72: 102–105.
- Brown, Jessica. 2010. "Knowledge and Assertion." *Philosophy and Phenomenological Research* 81: 549–566.
- Brown, Jessica and Cappelen, Herman (eds.). 2011. *Assertion: New Philosophical Essays*. Oxford: Oxford University Press.
- Davis, Steven. 1994. "Anti-Individualism and Speech Act Theory." In Savas Tsohatzidis (ed.), *Foundations of Speech Act Theory*. London: Routledge.
- DeRose, Keith. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.
- Gauker, Christopher. 2003. "Social Externalism and Linguistic Communication." In Jose Frapolli and Esther Romero (eds.), *Meaning, Basic Self-Knowledge, and Mind*. Stanford: CSLI Publications.

¹⁷We are grateful to John Hawthorne, Lisa Miracchi, Blake Roeber, Ernest Sosa, Kurt Sylvan, Angelo Turri, and to two anonymous referees for helpful feedback. This publication was made possible through the support of a grant from the John Templeton Foundation; the opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. This research was also supported by the Social Sciences and Humanities Research Council of Canada, the British Academy, the Association of Commonwealth Universities, the National Endowment for the Humanities, and an Ontario Early Researcher Award.

- Geach, P.T. 1965. "Assertion." *The Philosophical Review* 74: 449–65.
- Harnish, Robert. 2009. "Internalism and Externalism and Speech Act Theory." *Lodz Papers in Pragmatics* 5: 9–31.
- Kelp, Christoph. 2013. "A Practical Explication of the Knowledge Rule for Informative Speech Acts." *European Journal of Philosophy* 21: 367–383.
- Lackey, Jennifer. 2011. "Assertion and Isolated Second-Hand Knowledge." In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 251–275. Oxford: Oxford University Press.
- Searle, John R. 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.
- Turri, John. 2010. "Epistemic Invariantism and Speech Act Contextualism." *The Philosophical Review* 119: 77–95.
- . 2011. "The Express Knowledge Account of Assertion." *Australasian Journal of Philosophy* 89: 37–45.
- . 2013a. "Knowledge Guaranteed." *Noûs* 47: 602–612.
- . 2013b. "The Test of Truth: An Experimental Investigation of the Norm of Assertion." *Cognition* 129: 279–291.
- Weiner, Matthew. 2005. "Must We Know What We Say?" *The Philosophical Review* 114: 227–251.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.