

ON PREDICTING

FABRIZIO CARIANI

University of Maryland, College Park

I propose an account of the speech act of prediction that denies that the contents of prediction must be about the future and illuminates the relation between prediction and assertion. My account is a synthesis of two ideas: (i) that what is in the future in prediction is the time of discovery and (ii) that, as Benton and Turri recently argued, prediction is best characterized in terms of its constitutive norms.

PREDICTING is widely believed to be a matter of making assertions (or assertion-like speech acts) about the future. Both pieces of that understanding are questionable. Predicting is not essentially asserting, though it sometimes is; and the contents of acts of prediction are not essentially about the future, though they often are. I propose a theory of predicting that captures its relation to asserting and also captures the sense in which it appears to essentially involve the future.

My discussion is structured around two puzzles. The *subject-matter* puzzle (§1) challenges the thesis that the contents of predictions must be about the future. The *speech act* puzzle (§2) concerns the relationship between predicting and asserting. I lead with them to establish the goal posts of the inquiry. With the puzzles on the table, I identify the views of some fellow travelers who agree that predictions are not essentially about the future: I consider (§3) the view that predicting requires future discovery, and then (§4) a proposal by Benton and Turri (2014, henceforth *B&T*) to the effect that predictions are characterized by a constitutive norm (analogous to, but weaker than, the knowledge norm for assertion). I argue that each of these views captures something important about predicting and formulate an account that synthesizes their virtues (§5). It is distinctive of my proposal, and an important part of my argument, that the speech act of prediction is, in a sense, multiply realized: there are two different kinds of predictions, roughly lining up with those predictions that are also assertions, and those that are not.

Contact: Fabrizio Cariani <fabrizio@umd.edu>

The arguments to come require us to have minimally reliable heuristics for determining what counts as an act of prediction. I propose two, very rough, sufficient conditions: the first is that speech acts whose vehicle is a sentence of the form *I predict that ...* are predictions. Call these *performative predictions*. The idea behind this terminology is to analogize predictions to promises, apologies, and other speech acts that are canonically executed by performing another speech act—for example by asserting, or, as Searle and Vanderveken (1985) prefer, by *declaring* something. The classification of predictions as performatives is supported by standard tests for performativity such as “hereby”-insertion (Austin 1975), as in: *I hereby predict that they will lose the match*. And indeed, “predict” is standardly treated as a performative verb, for example in Searle and Vanderveken (1985). As a point of terminology, I say that the content of *I predict that p* is the proposition that *p*—as opposed to the proposition that the speaker predicts that *p*.

Evidently, not all predictions are performative. One can predict that it will snow overnight just by uttering the sentence *it will snow overnight*. More generally, one can predict that *p* by uttering a sentence that, in context, expresses the proposition that *p* (call these *transparent predictions*). How to recognize transparent predictions? I propose the *voiceover heuristic*: individual speech acts are classified as predictions if we can felicitously imagine a voiceover continuation like *that prediction turned out to be right (/wrong)*. Some examples:

- (1) A: this plane will land on time.
voiceover: that prediction turned out to be wrong.
- (2) B (staring out of a window): it’s raining.
voiceover: *that prediction turned out to be wrong.

There might be predictions that are neither performative nor transparent, but if so I set them aside.¹

One last piece of ground-clearing before we begin in earnest. For this paper, I set aside the theoretical meaning of “prediction”, as operative in claims like:

- (3) Special relativity predicts that a twin in a high-speed rocket, as viewed by her Earth-bound sister, will have a slower-ticking clock.²

There are clear limits to how tight the connection between predicting and theoretical predictions can be. To start, theoretical predictions differ from speech acts in their ontological status. Speech acts are widely understood to be events and

1. It is an established point in speech act theory that we should be wary of linking speech acts with hard-and-fast principles connecting a sentence’s form and the force of its uses (Green 2017). So there might well be predictions that are neither performative nor transparent. However, none of the points below strictly depends on this distinction being exhaustive, and ignoring this complication will streamline presentation.

2. <https://physics.aps.org/synopsis-for/10.1103/PhysRevLett.113.120405>

endowed with the kind of structure that is distinctive of acts. In particular, *qua* acts, they are events that have an agent. It is for this reason, after all, that they can meaningfully be said to be subject to norms—and in particular, to the sort of norms Benton and Turri (2014) and myself invoke in our respective characterizations. By contrast, theoretical predictions are not events, but propositions. In particular, the predictions of a theory *T* are some (but not all) of the propositions that *T* entails. As such, they have whatever ontology propositions have.³ If, for instance, propositions are abstract objects (e.g., sets of worlds), then theoretical predictions are those very same abstract objects. This suggests that if a normative theory of the speech act of prediction is correct, there is unified analysis that captures acts of prediction and theoretical predictions.⁴

1. Predicting ≠ Forecasting: An Easy Piece

Say that a forecast is a speech act whose content is entirely about the future. According to the *future content* hypothesis, predictions are forecasts in this sense.

Authors with disparate commitments and diverse backgrounds accept the future content hypothesis. According to Searle, prediction is an assertive (i.e., it signals commitment to truth and has word-to-world direction of fit) with future subject matter:

The differences [...] 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. (1985: 6)

Sperber and Wilson follow suit:

[W]hat makes an utterance a prediction is that [the speaker] ostensibly communicates an assumption with a certain property, that of being about a future event at least partly beyond her control. (1986: 245)

And here is Isaac Levi:

Prediction [...] can be understood to express full belief in the truth of some claim about the future or a judgment as to how probable some conjecture about the future is to be true. (2007: 1)

3. Of course, each act of prediction is also associated with a proposition—the content of the prediction. (Say that the contents of acts of prediction are *S*-predictions.) Unfortunately, the existence of *S*-predictions is not enough to forge a tight connections between theoretical predictions and acts of prediction. Although propositions are involved in both cases, *S*-predictions can, and theoretical predictions cannot, be characterized in terms of norms governing the behavior of an agent.

4. There is room, however, to theorize about how these concepts are connected. I have a proposal on offer in Cariani (in press).

The problem for the future content hypothesis is that some predictions are not about future events. It is possible to make predictions that are entirely about the present or the past (see also Benton & Turri 2014: §2). Here are two cases which illustrate two different ways in which predictions may concern past events.

- (4) Steph Curry is about to take a free throw. Elena is watching the game and says: “I predict he’ll make it”. But the game is actually on tape delay. Curry has already made the free throw.
- (5) Simona takes a test at 9 AM, then leaves for vacation. The test is graded instantly but Simona has not let her friends know the outcome. One of her friends goes on to say: “I predict that she passed”.

In the first case, the actual temporal reference seems to have been replaced by a “story-internal” temporal reference, as if the basketball game was a fiction endowed with its own temporal sequence. To illustrate, reading *Dr. Zhivago*, one might predict that Yuri Zhivago will die—meaning that he will die *by the end of the story*. Yuri’s death is not in the future with respect to the predictor’s context; it is in the future within the timeline of the story. Perhaps prediction can operate on the basis of such a simulated timeline. If that is true, then (4) does not by itself refute a suitably modified future content view.

Things are different with (5). The speech act in (5) meets our sufficient conditions for qualifying as a prediction: it is a flawless performative prediction.⁵ Moreover, (5) is a more striking counterexample than (4), since here there is no suggestion of futurity with respect to a surrogate timeline. What seems to have to be in the future in (5) is the time at which the participants will receive conclusive evidence on the question whether Simona passed (I consider an account that builds on this hunch in §3).

For another example in this vein, imagine a classroom of preschoolers playing a game in which they have to guess, without looking, what object is inside a box. The teacher might ask: *do you have any predictions about what is in the box?* A student answers: *it’s a toy car!* The student made a prediction but nothing about the subject matter of the prediction is in the future.

Our first puzzle then is: how should we conceive of prediction if not in terms of the future content hypothesis? The etymology suggests that predicting involves saying something *ahead of* something else. Of course, that etymology might well not be worth taking seriously. But if it is, we might wonder what that something else is, given that it’s not the content of the prediction that relates it to the future.

5. Note also that it is easy to imagine a simple variation on the story in (5) on which Simona’s friend makes a felicitous transparent prediction.

2. Predicting vs. Asserting

Onto the second puzzle: how are predicting and asserting related? Two conflicting ideas spring to mind.

the subkind thesis: every act of prediction is also an assertion.

the incompatibility thesis: no act of prediction is also an assertion.

A paradigmatic implementation of the subkind thesis is the idea that predictions are assertions about future states of affairs. A paradigmatic implementation of the incompatibility thesis is the idea that prediction and assertion are incompatible speech acts that share some traits because they belong to a common genus. An example of this view is the idea, mentioned above, that there is a class of “assertives” that includes speech acts as diverse as asserting and guessing. Of course, acknowledging that they have commonalities is consistent with the incompatibility thesis.

One’s choice among such options is consequential for the theory of assertion. Several theorists express the intuition that the epistemic standards for prediction are weaker than knowledge (e.g., Weiner 2005; Williams 1994). If that is accepted, and if one accepts the subkind thesis, the standard for assertion must sometimes be weaker than knowledge.⁶ Alternatively, a defender of the knowledge account of assertion might feel tempted to ride this argument the other way around, and conclude that predictions cannot be assertions.

The subkind thesis cannot be right. This is shown by focusing on an interesting and under-appreciated datum involving performative predictions. Unlike assertions, performative predictions can felicitously be followed by Moorean professions of ignorance in the predicted proposition.⁷

(6) I predict that she will win but I don’t know that she will.

This is important because it is a fundamental characteristic of assertions that they cannot be followed by professions of ignorance.

(7) * She won but I don’t know that she did.

6. There are views of assertion on which the standard for assertion is context-sensitive and can sometimes be knowledge, sometimes weaker than knowledge (Goldberg 2015; Levin 2008). In the main text, I have in mind a non-context sensitive knowledge norm.

7. A reviewer highlights, though not as an objection to this point, that it is bad to say:

(i) I predict that she will win but I don’t know that I predict that she will

This might be thought of as (mild) evidence that *I predict that A* is an assertion whose content is that I predict the proposition that A. Of course, it is important to keep in mind the gap between this and the further commitment of the subkind thesis, namely that a speech act whose vehicle is *I predict that A* predicts the proposition that A. It is this further commitment that is up for criticism in the main text.

Relatedly, performative predictions are permissible even if based on purely statistical evidence. In a fair lottery, one may felicitously utter *I predict that my ticket will not win*. But one may not outright assert *my ticket will not win*.

The upshot is that performative predictions fail to satisfy two important tests that are plausibly requirements of assertion. Such tests are often leveraged in support of the knowledge account of assertion (Williamson 2000). But accepting the tests as diagnostics for assertion does not require commitment to the knowledge account.⁸ All that is required is that assertions be incompatible with professions of ignorance, however that incompatibility might be explained, and that they be defective if made on the basis of purely statistical evidence. Given that sentences like (6) are felicitous, that they are predictions, and that their felicity does not exploit specific features of context, we must conclude that some predictions are not assertions.

It is no doubt possible to resist this argument by insisting that the tests are not requirements of assertion. The ensuing dispute would be partly terminological. Although I am not authorized to issue fines to those who delimit assertion differently than I do here, there is a theoretically valuable class of speech acts for which the tests *are* diagnostic. Referring to the class of speech acts that is pinned down by these tests as “assertions” is, in my view, a legitimate way to use the term. With that said, this argument can be extended by noting that there are related contrasts involving other concepts:

- (8) I predict that she will win but I am not committed to that (/I am not certain/she might not).
- (9) * She won but I am not committed to that (*I am not certain, */it might be that she did not).

These asymmetries are pervasive and distinguish performative predictions from ordinary assertions on a broad variety of views about assertion, including commitment views in the style of Brandom (1994), and context update views in the style of Stalnaker (1978).

The failure of the subkind thesis should not be taken to support the incompatibility thesis. In fact, I think the incompatibility thesis is also false: acts of transparent prediction have all the hallmarks of assertion. Of course, defenders of the incompatibility thesis would agree that there are important similarities between predictions and assertions, but still insist that they are different—recall the analogy with distinct species sharing a genus. However, if the incompatibility thesis were true, transparent predictions would have to lack some of the critical features of assertion. The problem for the incompatibility thesis is that they

8. Proponents of truth norms of assertion often agree with Williamson about such data and seek to account for them with different resources. See, for example, MacFarlane (2014: §5.5).

don't. For example, Benton (2012) argues that the key diagnostics that support the knowledge account of assertion also apply to *transparent* predictions. In particular, Moorean profession of ignorance is just as bad with transparent predictions as it is with past-directed assertions like (7).

(10) * She will win but I don't know that she will.

Weiner (2005) insists that predictive assertion is compatible with such Moorean denials. However, at the critical point in Weiner's argument in which Moorean data are discussed, it is clear that he has performative predictions in mind:

Suppose that after Aubrey asserts [that the French will attack at nightfall] Pullings asks, "How do you know that the French will attack at nightfall?" and Aubrey responds, *I don't know they'll attack at nightfall—we haven't intercepted their orders-but my prediction is that they will.* (Weiner 2005: 235, square brackets are my insertions)

Performative predictions are not probative in this context, since our question is precisely whether transparent predictions behave like them or in some other way. In sum, Benton's observation that (10) has the distinctive badness of Moorean denials seems strong and it is admitted into evidence as far as this paper goes.

This behavior of transparent predictions also sharply distinguishes them from "hedged assertions" (Benton & Van Elswyk 2020) as illustrated by:

(11) She won, I think, but I don't know for sure.

There is no reason to think that assertions about the future *have to be* hedged, and their incompatibility with Moorean denials helps emphasize that.⁹

Transparent predictions also behave like assertions in other important respects. As with assertions, it is defective to make transparent predictions only on the basis of statistical evidence. While Williamson (2000: 245) relies on past-directed examples (like *your ticket did not win*), it is just as bad—Weiner's (2005) intuition notwithstanding—to assert *your ticket will not win* on purely statistical evidence, absent any sort of hedge or special contextual cues to the effect that one is flouting

9. A reviewer identifies a much more promising way to appeal to hedged assertions. One might think that (i) *performative* predictions are hedged assertions and (ii) what I called "transparent predictions" are simply ordinary assertions. This proposal innovates over much of the literature I engage with here by obviating the need for a special theory of prediction. I have some credence in the first claim: while its development would be very different from what is to follow, there would be important connections between the resulting theories. I am much less concessive about the second claim. My discussion is founded on the strong intuition that *some* predictions involve no explicit hedges, nor any markers of performativity—not even ones that are inherited from the surrounding context. This judgment may not be universally accepted, but it is a sufficiently clear and strong judgment to warrant theorizing on it.

the norms of assertion.¹⁰

Taking a more general perspective, it would be bizarre if there was a conceptual barrier whereby simple declarative sentences such as *it will rain* just could not be used to make assertions.

I said that there would be a puzzle at the end of this discussion, and so far not much seems puzzling about these observations. Taken together, they guide us towards a surprising conclusion: performative predictions are not assertions and transparent predictions typically *are* assertions. Hence, there is no uniform relation between predictions and assertions: they are independent speech acts, in the sense that a single act can be both, one, or neither. I think this *is* the correct moral and, yes, not much of a puzzle, in light of the observations. But the observations are puzzling because not everyone will find this view congenial. Moreover, they do guide us towards an outstanding puzzle: how should we understand the speech act of prediction so that it's distinct from assertion but also such that it's clear that one and the same act can be *both* a prediction and an assertion?

3. The Future Discovery View

The key case I relied on in rejecting the future content hypothesis immediately suggests an alternative view. I repeat the case here with its original numbering:

- (5) Simona takes a test at 9 AM, then leaves for vacation. The test is graded instantly but Simona has not let her friends know the outcome. One of her friends goes on to say: "I predict that she passed".

The most distinctive feature of (5) is that, although the time of the event is in the past, the question whether Simona passed isn't settled until a time that follows the prediction. That fact appears to play an important role in explaining why the final utterance in (5) is a prediction.

We could stretch this insight into a view (*NB*: not *my* view). The *discovery view* agrees with the future content view that there is a future element in prediction. However, it diverges because it claims that what has to be future in a prediction that *p* is the time of discovery—that is, the time at which the question whether *p* is settled. Here is a slightly more precise statement.

The Future-Discovery Account of Predicting

A speech act *a* with content *p* is a prediction if and only if

10. Quick dialectic check: though I am using Weiner as a foil, he would agree that the incompatibility thesis is false and that transparent predictions are assertions. However, the reasons I am offering for that are strictly incompatible with Weiner's outlook. He doesn't take incompatibility with Moorean denials to be diagnostic of assertion.

(D0) *a* is an assertive.

(D1) the truth value of *p* is not settled by the evidence that is collectively available to the participants to the conversation.

(D2) the question whether *p* will be settled by forthcoming evidence.

It is difficult, perhaps impossible, to analyze the phrase “settled by evidence” as it occurs here. Still, without attempting to define it, we can try to illuminate it. Belief states are often viewed as closing off inquiry on some subject matter. According to this view, acquiring a belief that *p* is relevantly similar to adopting a plan: it involves some degree of close-mindedness about further deliberation on *p* (Friedman 2019). In the belief case, this close-mindedness is only appropriate when one’s evidence is strong enough to warrant closing deliberation. So, to make sense of what it is to be “settled by evidence”, we can think in terms of whatever level of evidence is strong enough for someone to acquire that warrant.

Setting aside the details of how this insight could be developed, it is evident that the discovery view classifies the speech act in (5) as a prediction. It has the features of assertives (commitment to truth, and word-to-world direction of fit), and it meets D1 and D2. The collective evidence leaves it open whether Simona passed the test, and it is reasonable in the provided context to take it for granted that the conversational participants will find out whether she did indeed pass. A similar treatment also applies to the preschool object-in-a-box example: when a child answers *it’s a toy car!*, she counts as making a prediction because the current evidence does not settle the contents of the box and because the class will soon discover what is in the box.¹¹ Surprisingly, the discovery view even accounts for the tape delay example in (4), without any additional assumptions involving surrogate timelines. What makes Elena’s speech act a prediction is that the time at which the relevant conversational participants will find out whether Steph made the freethrow is in the future.

Benton and Turri object that the discovery view—which they interpret somewhat differently than I do—attaches the wrong subject matter to predictions.

The primary problem with the [Discovery] 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. (*B&T*: 1861)

B&T are correct in emphasizing the importance of subject matter. However, not

11. This example also helps illuminate why condition D1 needs to refer to the collectively available evidence, as opposed to the individual evidence: suppose that one child already knows what object is in the box because she peeked ahead of time; when she says *it’s a toy car!* her utterance should still count as a prediction.

every version of the discovery view gets the subject matter of predictions wrong—and specifically the view in $D0-D1-D2$ doesn't. The versions of the discovery view that are affected by *B&T's* objection claim that *I predict that it will rain* is approximately synonymous with *I predict we'll discover that it will rain*. However, we do not have to spell out the discovery insight in this way. Future discovery might be a necessary condition for predicting, without it being reflected in the content of the speech act.

B&T advance a second, independent objection:

Late in life, Edwin Hubble lobbied for the Nobel 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. (*B&T*: 1861)

B&T's reasoning is that, according to the future discovery view, Hubble's speech act counts as a prediction only if Hubble got to find out whether the Nobel Prize reclassified astronomy. But he did not get to find out, so his speech act does not count as a prediction.¹²

This objection succeeds against versions of the discovery view that make it a necessary condition that the *speaker* will find out. But, here again, we must be careful to note that it does not strike against every version of the discovery view. The discovery view formulated above—though somewhat vague about what it is for evidence to be forthcoming—does not state this requirement. It only demands that the question be answered at some point relative to some body of evidence that is suitably related to the community in which the prediction takes place.

With that said, there is an immediate generalization of *B&T's* second objection that targets every version of the discovery view that includes a constraint like $D2$. Some predictions concern events that future evidence is not guaranteed to settle. Here is an example borrowing from an unrelated discussion in Dummett (1959):

(12) A city will never be built here.

If a city is built on the relevant spot, someone will indeed find out. But if no city is ever built on the relevant spot, no one will find out.¹³ Despite that, someone

12. *B&T* say that the prediction isn't "falsified", but I think it is not charitable to interpret them to mean this. Whether the prediction is true or false is not at issue here. What is at issue is whether the speech act is to be classified as a prediction. The reason *B&T* talk about the prediction being "falsified" is that they interpret the discovery view as claiming that *I predict that p* entails *I will discover that p*. I have already objected to that part of their view, but I think the present objection is instructive when reinterpreted along the lines that I do here.

13. The case might be spelled out so that someone can find out even if no city will never be

who uttered (12) ought to be counted as making a prediction—both by direct intuition and by the voiceover test.

The generalization in the neighborhood is that predicting that p does not require that it be settled that the question whether p will be conclusively answered at all. It does not have to be settled for the speaker; not for the addressee; and not for anyone else.

Another, related, reason for concern is that some predictions target conditional claims.¹⁴ Consider the statement:

(13) If we put cats on a spaceship, they would find a way to survive.

The discovery view fails to classify conditional predictions as predictions, and for the same reasons why it fails on potentially unsettled predictions like (12). Namely, it cannot be settled that the question concerning the content of the prediction will be answered. After all, conditional predictions weigh on conditional questions, and if the antecedent of the conditional question is not satisfied, then there is no expectation that the conditional question will be conclusively answered.

4. Normative Views

Benton and Turri leverage their criticisms of the future content and future discovery views into an alternative account. They maintain that the speech act of prediction is characterized by its constitutive norms—drawing on Williamson’s (2000) defense of the knowledge account of assertion. In particular, predictions are characterized by a norm that requires the speaker’s credence to have a particular kind of structure and a particular relation to the speaker’s evidence.

4.1. Formulating Normative Accounts of Prediction

Williamson characterizes assertion in terms of core rules (the “constitutive rules”), on the basis of an extended analogy between the rules governing speech acts and the rules governing games.¹⁵ Following and expanding on Unger (1975), built on the spot in question. What matters to my argument is that the case *can also* be set up so that no one will find out.

14. I say “predictions of conditional claims” instead of “conditional predictions” because the latter might be interpreted on an analogy with conditional assertion, which I do not intend to discuss here. A conditional assertion is an assertion that is made under a condition. If the condition is not satisfied, no assertion takes place (for a recent discussion of the theory of conditional assertion, see Goldstein 2019). By analogy, a conditional prediction ought to be a prediction that is made under some condition, so that the condition is not satisfied no prediction has taken place. It is an interesting question whether there are conditional predictions in this sense, but not one I will tackle here.

15. Williamson appropriates the phrase “constitutive rules” from speech act theory (e.g., Searle 1969), but emphasizes a critical respect that is underappreciated in the speech act theory

Williamson suggests characterizing assertion as the unique speech act whose constitutive rule is:

Knowledge account of assertion: one may assert p only if one knows p .

B&T follow a similar playbook in their normative account of prediction. According to them, prediction is the unique speech act subject to:

Proper expectation rule for prediction: one may predict p if and only if one properly expects p .

It is interesting and unusual that *B&T* state their norm in terms of a biconditional. It is standard protocol to state norms for speech acts as necessary conditions for permissibility—not as necessary *and sufficient* ones. The reason for the protocol is that if the force of the permission modal *may* is left unspecified, a performance of a speech act might be classified as impermissible for reasons that have nothing to do with the constitutive norms for that speech act. Having acknowledged this point, however, it's reasonably clear what *B&T* have in mind. They mean, as far as *constitutive* requirements go, proper expectation is the only standard: one may predict p only if one properly expects p ; any other form of impermissibility must result from violations of principles that are not constitutive of prediction. To see why this might work consider an analogy: the game of chess could be characterized in terms of a biconditional pinning down the set of permissible moves at any point in a game. Such a biconditional might not rule out rude behavior. But it would not follow from that that rude behavior is permitted during chess matches. The job of the biconditional is to pin down permitted behavior only as far *as far as the rules of chess go*.

Perhaps the most urgent clarification concerns what it is to “properly expect” something. According to *B&T*, the mental state of expectation is primitive. They gloss it as follows:

canon. For Searle, a constitutive rule r for practice π is such that if r is not obeyed then participants are not engaging in π . On this conception, constitutive rules both define a practice and state necessary and sufficient conditions for it. Williamson emphasizes that there is critical conceptual space between defining a practice and stating necessary and sufficient conditions for it.

Constitutive rules do not lay down necessary conditions for performing the constituted act. When one breaks a rule of a game, one does not thereby cease to be playing that game. When one breaks a rule of a language, one does not thereby cease to be speaking that language; speaking English ungrammatically is speaking English. Likewise, presumably, for a speech act: when one breaks a rule of assertion, one does not thereby fail to make an assertion. One is subject to criticism precisely because one has performed an act for which the rule is constitutive. (Williamson 2000: 240)

Williamson is right. The “travel” rule in basketball is part of a rule-based definition of the game of basketball. But if a player travels during a game, they are still playing basketball.

(mere) expectation is a mental state of slight commitment, which requires regarding a proposition as more likely than not. (*B&T*: 1862)

One expects *properly* when one's expectations are, in some sense, sanctioned as correct by one's evidence.

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*. (*B&T*: 1862)

It is particularly instructive to understand how *B&T*'s proper expectation norm relates prediction to assertion. As they highlight (*B&T*: 1865), the proper expectation account makes it possible for there to be overlap between predicting and asserting. At the same time, the proper expectation account also explains why not all predictions are assertions: predicting is a game that can be played for a cheaper price than asserting (i.e., by subjecting oneself to lower epistemic standards). In some contexts—when the information about the most urgent subject matters does not quite meet the epistemic standards for assertion—a speaker might wish to play the cheaper game only.

The proper expectation account threads nicely through our motivating puzzles. It does not demand that the contents of predictions be in the future, so it can account for the evidence that doomed the future content view. And it does not demand that future discovery be on the horizon. For that reason, it is immune to the criticisms we raised against the future discovery view. Someone who predicts that a city will never be built on this spot might well possess evidence that favors that claim over its negation.

4.2. *Evaluating the Proper Expectation Account*

Though the proper expectation account marks a striking advance in our understanding of prediction, there are some problems for it that make it desirable to look for a refinement.

Consider first what the proper expectation view predicts about present-directed certainties. Suppose I am staring out of my window and I see heavy rain. Obviously, I meet the conditions for properly expecting that it's raining. The credences that are appropriate to my evidential state all favor raining over the alternatives. But it does not seem plausible to say that I can predict that it's raining.¹⁶ The attempted performative prediction *I predict that it's raining* is infelicitous (or, in the right context, an attempt at a joke). I can, of course, assert

16. This case is reminiscent of a case discussed in von Fintel and Gillies (2010) to illustrate the evidentiality of epistemic *must*. Despite the different theoretical domain, there seems to be a connection between the requirement of "indirectness" that von Fintel and Gillies discuss for *must* and the parallel requirement that appears to be operative in prediction.

that it's raining; but in this case the voiceover heuristic suggests that I'm not, after all, predicting.

(14) A: It's raining.

voiceover: *that prediction turned out to be right.

B&T gesture towards an account of the infelicity of such predictions within their theory. They propose that there are strength relations between speech acts that are somewhat parallel to strength relations between contents. It is a basic Gricean point that (cooperative) assertion is governed by a "assert-the-stronger" rule. It can be defective (because uncooperative) to assert that one believes it's raining when one knows it's raining. Similarly, assertive speech acts might be governed by a "perform-the-stronger" rule.

If you're going to perform an assertive speech-act, you should perform the strongest assertive that your evidence permits (and no stronger).

So, the account goes, I cannot predict that it's raining because I can assert that it's raining. Because assertion is the stronger (i.e., more demanding) assertive speech act, that's what I should perform.

This reasoning fails for two independent reasons. To start, "perform-the-stronger" only operates as *B&T* expect it to when it applies to incompatible speech acts—speech acts that cannot be performed simultaneously. Part of what motivates the "assert-the-stronger" rule is that any one assertion can only be an assertion of *one* content. But it is a key claim of the present discussion that prediction and assertion are not incompatible: one and the same speech act can be both. Since they are governed by compatible constitutive rules, there is no metaphysical or normative barrier to performing an act which is subject to both sets of rules.

Perhaps *B&T* presuppose that, though they are governed by compatible constitutive rules, one can never perform both speech acts with a single event of utterance. But I don't see a reason to accept this. Two basketball teammates might engage in a competition about who can score more free throws with their off hand during a professional game. If so, their free throws would be simultaneously subject to both the rules of basketball and to the rules of their game-within-the-game. Indeed, this point is familiar from the theory of speech acts itself. A mainstream view of performative utterances, such as *I promise to clean my room* is that they involve two, overlapping speech acts: a promise to clean my room and another speech act such as an assertion, or a declaration, that I promise to do so (Searle & Vanderveken 1985: 4 and, more extensively, Ch. 7).¹⁷

17. A reviewer notes that an interesting theoretical perspective on this point emerges from the "semanticized" analysis of performatives in Condoravdi and Lauer (2011). Condoravdi and

A second problem with *B&T*'s account of why one cannot predict past and present certainties is that, much like the Gricean rules on which it's mirrored, "perform-the-stronger" requires a cooperative context. As Mitchell Green (1995) notes, Grice's Maxim of Quantity seems suspended in certain contexts of strategic conversation—for example, if I am testifying in a court setting.¹⁸ While the strategic context still requires a substantial amount of cooperation, there need to be no presumption that one be informative to the point that one must always assert the stronger of two (relevant) propositions. The fact some inferences traditionally associated with the maxim of quantity still seem to go through in such contexts is evidence that they aren't entirely based on the kinds of reasoning outlined by Grice.

Something similar should also hold of "perform-the-stronger". Since the rule is based on the insight that we should maximize the informative upshot of our speech acts, it ought to be suspended in contexts that, for whatever reason, are less than fully cooperative. And yet it is still bad in such contexts to make predictions about things that one is certain of. Even in the courtroom, I cannot predict *p* when I and my interlocutors share direct, conclusive evidence in favor of *p*.

Reflecting on such cases leads to another concern for the proper expectation account. Though I have argued at length (and concur with *B&T*) that there can be predictions about the past, their distribution is limited in ways that the proper expectation account does not address. One can predict past and present eventualities only if there is some reason to believe that one will find out what that answer is. Suppose that, a few decades out of high school, I am looking at the photo of a basketball team I used to play on back then. The photo reminds me of my teammate Sam—who used to be a passionate, committed player and a remarkable athlete. It is odd for me to say:

(15) *? I predict that Sam still plays basketball to this day.

The reason why this is odd need not be that I don't have a proper expectation. I may have excellent evidence that people with Sam's skill, passion, and commitment tend to stay active as players into their middle-age.

What I suggest is that the extent to which (15) is licensed depends on the degree to which we have reason to expect direct evidence for it. This is the kind of phenomenon that the discovery view gets right. If we lack a reason to

Lauer make a strong case that performative utterances are in the first instance assertions. On my reading, however, Condoravdi and Lauer are generally neutral on the ontological question whether there are multiple speech acts associated with a performative.

18. See also Asher and Lascarides (2013). Furthermore, recent work in philosophy of language, such as Camp (2018), draws attention to the fact that non-cooperative contexts are probably the norm, and not outliers to be excluded from the analysis.

expect that we will find out, (15) sounds bad. However, if we imagine modifying the context, so as to suggest that we will (or might) find out, the status of the prediction improves. So, for instance, suppose I am about to attend the twentieth year reunion of the basketball team; I expect to learn about what Sam ended up doing. In such a context, (15) may well be a flawless utterance.

A third concern, which in the context of this paper I will leave somewhat vague, is that the proper expectation view says too little about the relation between predicting and theoretical prediction. In (Cariani in press: §9.7), I explore the idea that if we accept the synthetic view, there are some links one can draw between theoretical prediction and the speech act of prediction. That account of the connection between these does not generalize plausibly to the proper expectation account.

As a last remark about the proper expectation view, I note a point of detail. *B&T* say that an expectation is proper when the evidence supports p over $\neg p$. This is too strong. I can non-defectively predict that Spain will win the world cup without being willing to take them over the field. The norm does not capture this, since it requires my evidence to support the proposition that Spain will win more strongly than its negation. The norm ought to be weakened so as to say that one can predict that p as long as evidence supports having greater credence in p than in the relevant alternatives. This observation is not an objection to the proper expectation account, since it can be addressed with an easy tweak. Instead, it is best taken as a reason to redefine what counts as a proper expectation. Towards that goal, say that one *properly expects* that p when (i) one's credence in p is greater than one's credence in the alternatives and (ii) one's evidence supports these inequalities. In the remainder, I will use *proper expectation* in this more generous sense.

5. The Synthetic View

I propose an account of predicting that synthesizes elements of both the discovery view and the proper expectation view.

5.1. Introducing the Synthetic View

Synthetic view. Prediction is the unique speech act \mathcal{A} that is governed by the rule that one may perform \mathcal{A} with content p in context c only if:

(s1) The contextual evidence in c does not settle whether p .

(s2) It is possible, as far as contextual evidence goes, that evidence that settles the question whether p is forthcoming.

(s3) One properly expects p .

Both conditions s_1 and s_2 appeal to a notion of “contextual evidence”. This could be taken to be the common knowledge of the participants to the conversation. Alternatively, in a more flexible sense, it might be taken to be some body of evidence that is salient in context. Under both conceptions, the contextual evidence is not solely determined by the speaker’s own epistemic state. In fact, under the flexible interpretation, there might be no connection at all between the contextual evidence and the speaker’s evidence—roughly as in “exocentric” uses of epistemic modals (Egan, Hawthorne, & Weatherson 2005; Stephenson 2007). I will entertain both conceptions in the following, though I will use the common knowledge interpretation as a default.

One point that deserves highlighting is that the sense of possibility in s_2 should be a “thick” one. For every unsettled question Q , there is a smidgen of epistemic possibility that Q will be settled in the future. So, it cannot be that any degree of possibility is enough to meet s_2 . That would trivialize s_2 , since it would follow that whenever s_1 is satisfied, so is s_2 . Since we are already demanding satisfaction of s_1 , s_2 would be idle. For s_2 to not be idle, the relevant sense of *possible* should be closer to *seriously possible*. I won’t attempt to analyze this concept of serious possibility, relying instead on general linguistic competence to evaluate individual applications of the theory.

It is possible to weaken clause s_3 to the claim that one simply expects p , by dropping the propriety condition. There is a strong intuition here, which s_3 accounts for, that one ought not to predict against one’s evidence. However, that intuition might be accounted for by claiming that propriety is a normative requirement on the mental state of expectation. It won’t matter to my analysis whether we think of propriety as built into the analysis of the speech act of predicting, or as a separate requirement on mental states of expectation.

5.2. Evaluating the Synthetic View

The synthetic view threads with ease through the desiderata we accumulated up to this point. Recall from §3 that we took *B&T*’s first objection to demand that a discovery-based account of predicting should separate the discovery component from the content of the prediction. At that point, I noted that that can be accomplished by letting future-discovery be a *necessary* condition of predicting and keeping claims about discovery out of the content of prediction. The synthetic proposal also implements something similar to this blueprint, but in a normative setting. According to the synthetic proposal, *normatively* correct predictions (i.e., predictions that comply with the constitutive rules) require the possibility of future discovery. This means that we are twice removed from the original proposal: the synthetic view requires less (serious possibility of future discovery vs. guaranteed future discovery). Moreover, the sense of “requires” is different. For the original discovery view, future discovery was a necessary condition of prediction.

In the context of the synthetic view, the relevant requirement is normative: the possibility of future discovery is a requirement of *permissible* prediction.

The synthetic view also deals handily with *B&T's* Hubble example. Future discovery need not be up to the speaker. All that has to be the case for something to count as a prediction is that future discovery be possible in light of the contextual evidence. In §3 I noted some related challenges to the Hubble example. For instance, I noted that the naïve future discovery view fails in cases like predictive utterances of:

(16) A city will never be built here.

The synthetic view handles this by merely requiring that evidence settling *p* might become available. In the city case, we know that if the prediction is true we will never find out. However, if the prediction is false, it is possible that the target epistemic state will be informed about it. Very much the same treatment can be extended to predictions of conditional contents. Suppose I predict that if Liz runs, she will win. It is typically not the case that we are guaranteed to find out whether this prediction is correct. But all that the hybrid view requires is that this be possible. And indeed it is possible unless we are certain that Liz doesn't run, or certain that if Liz runs she won't win (in both of these exceptional cases, the prediction would intuitively be defective).

Discussion of such cases naturally draws attention to the status of those predictions that occur in contexts in which it is settled that we will never find out either way. Suppose, as a matter of fact, we will never find out either way whether the number of stars is infinite and that it is common knowledge that we won't. The synthetic view predicts that one cannot meet the norm when one predicts that the number of stars is finite. A speaker might well be making a prediction, but they would be making a defective one. I take this to be the correct characterization of the case. It does seem that it is in some sense defective to make such a prediction, even if one had some evidence for it and if one apportioned one's credence to the evidence in the correct way. The synthetic proposal explains this defectiveness without relying on any additional norms. Those who disagree with me on this verdict, might still accept a version of the synthetic view that scratches requirement *s2*.

Let us move on to the constraints I used to critique *B&T's* account. First, I mentioned the case of present-directed certainties that are settled by one's direct evidence. I considered a case in which I am staring out of a window at a rainstorm and noted that in such case I cannot felicitously predict that it's raining. The synthetic view explains this as a violation of *s1*. Second, in the case of the high school basketball photo, the prediction plausibly violates clause *s2*. In the initial version of the case, we are not given any reason to expect that we will find out. In the variant in which we imagine that we are about to attend a reunion, we can

feliculously perform the prediction. After all, it is now possible (indeed, likely) that we will find out whether Sam still plays basketball. In general, the synthetic view predicts, correctly in my view, that the appropriateness of predicting will covary with whether there is enough reason to think that we will discover.

Let us stress test the synthetic view one more time. What does the view predict about cases in which an agent wrongly believes that they violate the standards for prediction? Consider a small variant on Simona's exam case; Simona's friend Louise wrongly believes that her evidence settles that Simona passed. As it happens, that is not true: Louise's evidence does not settle that Simona passed. Louise says: *I predict she passed*. This utterance seems defective. But did Louise violate a norm of prediction? It would seem not: the synthetic norm is entirely insensitive to Louise's beliefs on the matter. Here I think we should fall back on the idea of "secondary propriety" that is sometimes discussed in connection with norms of assertion.¹⁹ The idea is that our behavior (generally speaking) is governed by a meta-norm according to which, if an agent is subject to rule *R*, then the agent ought to act in such a way that they know, or at least have enough evidence that, they do not violate *R*. For example, in a high-stakes basketball game, it might be irresponsible to run too close to the side line, even if one is actually inbound. As far as the rules of basketball go, of course, the only thing that matters is whether the agent is inbound or out of bounds. But as far as our overall evaluation of the agent goes, her walking too close to the line might be defective—perhaps even deserving of blame. In a similar way, it might be irresponsible, and thus defective, when one does not know that one is complying with the norm, and indeed Louise does not she is complying with the norm.

As a final point of evaluation, I do note that the synthetic view may be somewhat weakened to suit alternative theoretical frameworks. In recent papers, Mona Simion and Christoph Kelp have challenged Williamson's "constitutive rules" approach to speech acts.²⁰ According to Simion and Kelp, if some activity is governed by a set of constitutive rules, then an agent cannot systematically violate a (significant) majority of those rules without being seen as ceasing to engage in that activity. But, they argue, one can violate the knowledge norm systematically and still engage in the practice of assertion. As examples, they consider a thinker who under the spell of an evil demon is made to be such that the near totality of her beliefs are mistaken. This hapless victim might still assert things. However, their assertions would be systematically false. I have some doubts about the success of this challenge,²¹ but the ecumenical point I want to

19. See Williamson (2000: 256–257), DeRose (2002), Weiner (2005: 236). Lackey (2007) argues against the idea of secondary propriety. I find Lackey's arguments striking but not inescapable, though I won't face the escaping challenge here.

20. Kelp and Simion (2020) and Simion and Kelp (in press). These papers draw on, and aim to strengthen, related arguments by Maitra (2011).

21. I think that the agent we imagine when we try to imagine someone whose beliefs are

stress is that one can grant the challenge and still tell a story about constitutive rules in a lighter sense. This is exactly what Simion and Kelp do to make sense of the connection between knowledge and assertion. They propose that the generation and distribution of knowledge is an essential part of the *function* of assertion. In this sense, it might be seen as “constitutive” in a lighter sense: it is essential to refer to knowledge in a specification of the good functioning of an asserting agent. This lighter sense of constitutivity is also a plausible fallback for the synthetic view of prediction.

5.3. Predicting and Asserting

Where does the synthetic view leave us when it comes to the relationship between prediction and assertion? Suppose, for definiteness, that assertion is governed by the knowledge rule. Then, prediction and assertion are governed by compatible norms, so that it is metaphysically and normatively possible for a single act *a* to be both. *A fortiori*, the same of course is true for norms that set standards for assertion that are weaker than knowledge.²²

It is critical to this result that, even though someone who asserts *p* must know the content of their prediction, *S*'s knowing that *p* is compatible with there being contextual uncertainty about *p*—uncertainty of the sort that is required to license prediction. For example, if “unsettled by contextual evidence” is understood in terms of compatibility with common knowledge, then all that is necessary for the norms to be simultaneously satisfied is that (i) the speaker knows *p* (ii) but it is not common knowledge that *p* among the conversational participants and (iii) it is compatible with common knowledge that an answer to the question whether *p* is forthcoming.

At the same time, the relative weakness of the synthetic norm explains why there can be predictions that are not assertions (and in this respect it behaves in exactly the same way as *B&T*'s norm). In some contexts, the best one can do is predict without asserting.

Acknowledgments

For correspondences, conversation, and feedback on this material, Sandy Goldberg, Simon Goldstein, Megan Hyska, Christoph Kelp, Carlotta Pavese, Geoff Pynn, Paolo Santorio, Mack Sullivan, Malte Willer, audiences at King's College

nearly totally false is in fact not someone whose beliefs are nearly totally false. And if it was, very ordinary conversation with them would sound like conversations with bad chatbots.

22. This applies to all the weaker norms we have considered, specifically, the truth norm advocated by Weiner (2005), or the reasonable belief norm of Douven (2006) and Lackey (2007), or the context-sensitive norm of Levin (2008) and Goldberg (2015).

London, and at Northern Illinois University, as well as my Fall 2019 graduate seminar. Also special thanks to two anonymous referees for *Ergo* whose comments have greatly improved the paper.

References

- Asher, Nicholas and Alex Lascarides (2013). Strategic Conversation. *Semantics and Pragmatics*, 6(2), 1–62. <https://doi.org/10.3765/sp.6.2>
- Austin, John Langshaw (1975). *How to Do Things with Words*. Oxford University Press.
- Benton, Matthew A. (2012). Assertion, Knowledge and Predictions. *Analysis*, 72(1), 102–105.
- Benton, Matthew A. and John Turri (2014). Iffy Predictions and Proper Expectations. *Synthese*, 191(8), 1857–1866. <https://doi.org/10.1007/s11229-013-0377-y>
- Benton, Matthew A. and Peter Van Elswyk (2020). Hedged Assertion. In Sanford Goldberg (Ed.), *The Oxford Handbook of Assertion*, 245–263. Oxford University Press.
- Brandom, Robert (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.
- Camp, Elisabeth (2018). Insinuation, Common Ground and the Conversational Record. In Daniel Fogal, Daniel W. Harris, and Matt Moss (Eds.), *New Work on Speech Acts* (40–66). Oxford University Press. <https://doi.org/10.1093/oso/9780198738831.003.0002>
- Cariani, Fabrizio (in press). *The Modal Future*. Cambridge University Press.
- Condoravdi, Cleo and Sven Lauer (2011). Performative Verbs and Performative Acts. *Proceedings of Sinn und Bedeutung*, 15, 149–164.
- DeRose, Keith (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2), 167–203. <https://doi.org/10.1215/00318108-111-2-167>
- Douven, Igor (2006). Assertion, Knowledge, and Rational Credibility. *The Philosophical Review*, 115(4), 449–485. <https://doi.org/10.1215/00318108-2006-010>
- Dummett, Michael (1959). Truth. *Proceedings of the Aristotelian Society*, 59(1), 141–62. <https://doi.org/10.1093/aristotelian/59.1.141>
- Egan, Andy, John Hawthorne, and Brian Weatherson (2005). Epistemic Modals in Context. In Gerhard Preyer and Georg Peter (Eds.), *Contextualism in Philosophy* (131–170). Oxford University Press.
- von Fintel, Kai and Anthony Gillies (2010). Must . . . Stay . . . Strong! *Natural Language Semantics*, 18(4), 351–383. <https://doi.org/10.1007/s11050-010-9058-2>
- Friedman, Jane (2019). Inquiry and Belief. *Noûs*, 53(2), 296–315. <https://doi.org/10.1111/nous.12222>
- Goldberg, Sanford C. (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press. <https://doi.org/10.1093/acprof:>

- oso/9780198732488.001.0001
- Goldstein, Simon (2019). The Theory of Conditional Assertion. *The Journal of Philosophy*, 116 (6), 293–318. <https://doi.org/10.5840/jphil2019116620>
- Green, Mitchell (1995). Quantity, Volubility and Some Varieties of Discourse. *Linguistics and Philosophy*, 18(1), 83–112. <https://doi.org/10.1007/BF00984962>
- Green, Mitchell (2017). Speech Acts. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Retrieved from <https://plato.stanford.edu/archives/win2017/entries/speech-acts/>. <https://doi.org/10.1093/acrefore/9780199384655.013.200>
- Kelp, Christoph and Mona Simion (2020). The C Account of Assertion: A Negative Result. *Synthese*, 197(1), 125–137. <https://doi.org/10.1007/s11229-018-1760-5>
- Lackey, Jennifer (2007). Norms of Assertion. *Noûs*, 41(4), 594–626. <https://doi.org/10.1111/j.1468-0068.2007.00664.x>
- Levi, Isaac (2007). Deliberation Does Crowd Out Prediction. In T. Ronnow-Rasmussen, B. Petersson B., J. Joseffson, and D. Egonsson (Eds.), *Homage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. Retrieved from <https://pdfs.semanticscholar.org/5770/ob7a82440cd93a7c5e7b3d7coa8fcde7c72f.pdf>.
- Levin, Janet (2008). Assertion, Practical Reason, and Pragmatic Theories of Knowledge. *Philosophy and Phenomenological Research*, 76(2), 359–384. <https://doi.org/10.1111/j.1933-1592.2007.00136.x>
- MacFarlane, John (2014). *Assessment Sensitivity*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199573004.003.0012>
- Maitra, Ishani (2011). Assertion, Norms, and Games. In Jessica Brown and Herman Cappelen (Eds.), *Assertion: New Philosophical Essays* (277–296). Oxford University Press.
- Searle, John R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139173438>
- Searle, John R. (1985). *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge University Press.
- Searle, John R. and Daniel Vanderveken (1985). *Foundations of Illocutionary Logic*. Cambridge University Press. https://doi.org/10.1007/1-4020-3167-X_5
- Simion, Mona and Christoph Kelp (in press). The Constitutive Norm View of Assertion. In Sandy Goldberg (Ed.), *The Oxford Handbook of Assertion*. Oxford University Press.
- Sperber, Dan and Deirdre Wilson (1986). *Relevance: Communication and Cognition*. Harvard University Press.
- Stalnaker, Robert (1978). Assertion. In *Context and Content* (78–95). Oxford University Press. https://doi.org/10.1163/9789004368873_013
- Stephenson, Tamina (2007). Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4), 487–525. <https://doi.org/>

10.1007/s10988-008-9023-4

Unger, Peter K. (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.
<https://doi.org/10.1215/00318108-114-2-227>

Weiner, Matt (2005). Must We Know What We Say? *The Philosophical Review*,
114(2), 227–251.

Williams, John N. (1994). Moorean Absurdity and the Intentional 'Structure' of
Assertion. *Analysis*, 54(3), 160. <https://doi.org/10.1093/analys/54.3.160>

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