Collective Inquiry and Expiring Assertibility

Anhui Huang

(Draft Oct 2026; please do not cite or circulate)

1 Introduction

On a natural view, whether an agent can felicitously assert a proposition at a time depends primarily on the strength of their evidence. It is puzzling, then, that in many cases an agent can felicitously assert a proposition at an earlier time but cannot felicitously assert that same proposition at a later time, *despite having the exact same evidence*. I will call this puzzle, introduced by Ninan (2022), the *expiring assertibility puzzle*. Consider (Ninan 2022, 406-7):

> **Rain**
> On Friday, Sam watched the weather forecast, which predicts that it will rain on Saturday. Later that day, his friend Julie asks him whether the amusement park will be open tomorrow. Sam says:
>
> (1) Probably not. It will rain tomorrow.
>
> On Sunday, another friend, Emma, asks Sam if the amusement park was open on Saturday. Without checking the weather reports or learning any new information, Sam says:
>
> (2) Probably not. #It rained yesterday.

Sam's utterance of (1) is okay, but his utterance of (2) sounds infelicitous. In uttering (2), Sam seems to be misleadingly portraying himself as having more evidence than he does. Instead, a hedged version of (2) would sound much better:

(3) It probably rained yesterday.
(4) I think it rained yesterday.

Sam's evidence remains the same throughout. This evidence is what put Sam in a position to assert (1). After the fact, however, without losing evidence, Sam loses the standing to assert the same proposition. This is puzzling: how can one lose the standing to assert when one's beliefs and evidence remain the same? It is as if the assertibility of the proposition simply expires. But why?

Intuitively, it has something to do with the kind of evidence Sam does and doesn't have: what he has—the weather forecast from Friday—just does not seem good enough for asserting (2). Ideally, a solution to the puzzle should sharpen and vindicate this intuition.

Solving the expiring assertibility puzzle has turned out to be difficult. In fact, it is difficult simply to articulate a clear generalization for when the puzzle arises: there are cases which are structurally identical to **Rain** in which assertibility does not expire. These are instances of *stable assertibility*. Consider (Ninan 2022, 433):

Vacation
On Thursday, Jack tells Alice that he is going to New York for the weekend. Later that day, a mutual friend, Jill, asks how Jack is doing, and Alice says:

(5) He is great! He is going to spend the weekend in New York.

Next Monday, Alice runs into a different mutual friend, Jane, who asks how Jack is doing. Alice says:

(6) He is great! He just spent the weekend in New York.

This time, both (5) and (6) seem fine. What distinguishes **Vacation** and **Rain,** such that one's standing to assert expires in one but not the other?

In this paper, I propose a novel solution to the expiring assertibility puzzle: Sam's assertion of (2) is inappropriate because it violates the *pertinent evidence norm of assertion*—you should assert a proposition only if you have pertinent evidence for it. Roughly, my view says that assertibility expires when the agent's evidence is no longer the best normally available evidence. Why? Because such evidence is needed to appropriately close inquiry, and an assertion is a proposal to close collective inquiry. In §2, I explain three existing kinds of solutions, the epistemic (Ninan, 2022; Torre, 2024), the modal (Cariani 2021), and the illocutionary solutions, and the problems they each face. In §3, I propose my own solution. I first argue that there is a pertinent evidence norm of assertion, and then show that it helps explain both when and why assertibility expires. I conclude in §4.

2 Existing solutions: epistemic, modal, and illocutionary

There are three kinds of solutions: epistemic (Ninan, 2022; Torre, 2024), modal (Cariani 2021), and illocutionary.[1] According to the epistemic solution, what it takes to make a past-directed assertion and a future-directed assertion is the same: knowledge of the proposition asserted.[2] What explains the asymmetry between (1) and (2) is Sam's epistemic position: he loses knowledge and therefore the standing to assert. According to the modal solution, Sam's epistemic position remains the same throughout. However, it takes less to make a future-directed assertion because "will" is a modal, which suppresses the standard evidential requirements. According to the illocutionary solution, the entire puzzle is based on a misunderstanding that in uttering (1) and (2), Sam is making the same speech act, i.e., assertion. Instead, Sam's utterance of (1) is a different kind of speech act, *predicting*, and predicting is subject to much weaker constraints.

2.1 The epistemic solution

The epistemic solution, endorsed by Ninan (2022) and Torre (2024), takes two views on board as a starting point. First, it assumes the knowledge norm of assertion (KNA):

(KNA) One should assert that p only if one knows that p.

[1] There isn't any published work that endorses the illocutionary solution. However, it has been discussed in Ninan (2022), which Ninan credited to an anonymous referee.
[2] Present-directed assertions fall on the same side of the asymmetry as past-directed assertions. For simplicity, I will only mention past-directed assertions.

Secondly, it follows Lewis (1996) (and others) in accepting a relevant alternative account of knowledge. Roughly:

> (RA) One knows that p iff p is true in every possibility which is (a) *relevant* and (b) uneliminated by one's evidence.[3]

The final and key step, then, is to argue that *time* affects what counts as relevant, or more precisely:

> (TAR) A possibility (possible world) can be irrelevant for an agent at time t and relevant for that same agent at t′ while the agent's evidence remains the same.

Ninan (2022) argues for TAR by appealing to the notion of normality. Intuitively, given a past and present, the way the future unfolds can be more or less normal. For example, given how the world has been up until now, a future in which it snows in San Francisco next winter is less normal than a future in which it does not snow in San Francisco. However, suppose it does snow in San Francisco next winter. Then given that it snows, a future in which people go swimming in San Francisco is less normal than a future in which people have a snowball fight. Relying on this intuitive notion of normality, Ninan proposes the following rule:

> (Rule of future normality) All else being equal, a world w′ is relevant at time t in world w iff: (a) w′ is *prima facie* relevant (b) w′ is at least as normal as w (given everything that has happened up until t).[4]

In other words, at any given time, we are allowed to ignore worlds in which the future unfolds in an abnormal way compared to the actual future. This offers us an explanation of why assertibility expires in **Rain**: Sam knows that it will rain beforehand because he can ignore *future abnormalities*, and he loses that knowledge afterwards because he cannot ignore *past abnormalities*.

Suppose that Sam's evidence at t_1 (when Sam utters (1)) eliminates all but one world in which it does not rain on Saturday: one in which a sudden drop in temperature after t_1 turns the rain into snow. Call this world **Snow**. Thanks to the rule of future normality, Sam knows that it will rain at t_1 because, given how the world has been up to t_1, **Snow** is less normal than the actual future and is therefore irrelevant for Sam at t_1. However, from the perspective of t_2 (when Sam utters (2)), **Snow** is not less normal than the actual world, given everything that happens up to t_2. After all, it is fairly normal for cold temperatures to turn rain into snow. So **Snow** becomes relevant for Sam at t_2. And since **Snow** remains uneliminated by Sam's evidence, Sam no longer knows that it rained on Saturday. Combined with KNA, this explains why Sam's assertion of (1) is okay whereas his assertion of (2) is infelicitous.

Torre (2024) provides a different argument for TAR. According to Torre, there is a fixed set of belief-forming processes that we generally take to be reliable.[5] Given a particular

[3] Here is a longer, more precise version:
> (RA′) An agent x knows proposition p in world w at time t iff p is true in every possibility w′ such that (a) w′ is relevant for x in w at t, and (b) w′ is uneliminated by x's evidence in w at t.

[4] Condition (a) is supposed to include all other considerations that affect whether a possibility is relevant while remaining neutral about what those are.

[5] They include processes like testimony and induction, for example, and exclude processes like clairvoyance and guessing.

conversational context and an assertion, there are presuppositions about which subset of these reliable belief-forming processes are in play for the proposition asserted. Furthermore, we are entitled to take this subset of reliable belief-forming processes for granted: we can properly ignore any possibility in which these processes produce false results.

> (Rule of reliable belief-forming process) All else being equal, a possibility w′ is relevant in a context c iff: (a) w′ is *prima facie* relevant (b) the reliable belief-forming processes presupposed in c produce true results in w′.

Torre's proposal offers another explanation of the expiring assertibility puzzle. Typically, for past-directed assertions, we presuppose that the reliable belief-forming processes at work are perception, memory, or testimony; whereas for future-directed assertions, we presuppose that they are induction, mental simulation, or intention formation. This is the case in **Rain**. When Sam utters (1), it is presupposed that the relevant belief-forming process is induction. As a result, Sam can ignore possibilities in which induction produces false results, thereby ignoring possibilities in which it does not rain on Saturday. However, when Sam utters (2), it is presupposed that the relevant belief-forming process is something causally downstream, such as testimony or perception. Now Sam is no longer entitled to ignore possibilities in which induction produces false results, so he loses the knowledge that it rained on Saturday.

The core idea of both Ninan's and Torre's proposals is that some possibilities can be ignored in the past but not in the future: for Ninan, these are possibilities where the future is abnormal; and for Torre, these are possibilities where induction (or mental simulation, intention formation) is unreliable. This makes it possible for an agent to lose knowledge by simply passing through time. However, there is a problem for both of their explanations, and for any epistemic solution to the expiring assertibility puzzle. The epistemic solution depends essentially on denying that Sam knows that it rained on Saturday when uttering (2). As a result, it struggles to explain some other linguistic data. For example, Cariani (2021) points out that, in cases of expiring assertibility, a kind of knowledge attribution seems true:[6]

(7) Sam knows that it must have rained on Saturday.

If we accept (i) that must-p entails p, and (ii) that knowledge is closed under entailment, both of which seem quite plausible, then it follows from the truth of (7) that Sam knows that it rained on Saturday. In addition, although the bare assertion is infelicitous, it seems that Sam *can* felicitously claim that he has knowledge by uttering (8):

(8) I know that it rained, because the weather forecast said so and it has never been wrong before.

The epistemic solution cannot explain why both (7) and (8) are felicitous.

[6] Ninan (2022) anticipates a similar observation: it seems appropriate for Sam to assert "It *must* have rained." However, he eventually denies that Sam is strictly speaking in a position to assert that. Instead, he *would* be in a position to assert that it must have rained, if he were to actually utter it. This response does not threaten (7), however, since it is a straightforward knowledge attribution, rather than an assertion.

2.2 The modal solution

In making the future-directed assertion (1), Sam uses the term "will." There are two views on the meaning of "will." On the tense view, "will" is a tense operator that shifts the time of evaluation forward. For example, "It will rain" is true at time t in world w iff there is some later time t′ in w such that it rains at t′ in w.[7] On the modal view, "will" is a modal operator, similar to "might" or "must." Cariani (2021) argues for the modal view and suggests that, by adopting the modal view, we can solve the expiring assertibility puzzle.[8] Unlike standard analyses of modals, which treat modals as quantifiers over possible worlds, Cariani treats "will" as a selection function. This selection function takes the world of evaluation as input and selects a world from its *historical alternatives*, i.e., worlds that are identical to it up until the time of evaluation. This set of worlds is the *modal base* of "will." The selection function is subject to exactly two constraints:

(Success) If the modal base is non-empty, it selects a world from the modal base.
(Centering) If the world of evaluation is in the modal base, it selects it.

Then, will-p is true just in case p is true at the world selected by the selection function.[9]

Typically, according to Cariani, an assertion carries an evidential requirement that the speaker has *causally downstream* evidence. However, modals cancel evidential requirements. So when Sam utters (1), the "will" in his sentence blocks the evidential requirement that he would have to satisfy otherwise; but when Sam utters (2), the evidential requirement goes through, and his failure to meet the requirement makes his assertion infelicitous.

Regardless of whether the modal view is a plausible account of "will", it faces a problem when used to solve the expiring assertibility puzzle, as Cariani himself acknowledges. For Cariani, standard assertions always carry the evidential requirement for causally downstream evidence; it is only because future-directed assertions involve modals that this requirement is canceled. This explanation overgeneralizes: it falsely predicts that assertibility would expire in every case structurally similar to **Rain**, including paradigm cases of stable assertibility like **Vacation**. More generally, it is just not true that asserting a sentence without modals *always* requires causally downstream evidence. For example, one can normally assert (9) and (10) without any causally downstream evidence.[10]

(9) My brother did not adopt a bunch of gorillas yesterday.
(10) The sun rose from the east in Monaco yesterday.

Cariani's modal solution falsely predicts that (9) and (10) are infelicitous unless the speaker has causally downstream evidence.

---

[7] This is also known as "Ockhamist semantics." Context might supply additional constraints for t′.
[8] What follows is an extremely brief and rough summary of Cariani's modal semantics. For the full account, see Cariani and Santorio (2018) and Cariani (2021).
[9] This is a rough and imprecise summary of the semantics. Here are the actual semantics:
(selection function) A selection function ***sel*** over a set of worlds is a function from $P(\Omega) \times \Omega \mapsto \Omega$ subject to two constraints.
(success) For all $\mathbf{A} \in P(\Omega)$ and all $w \in \Omega$, if $\mathbf{A} \neq \varnothing$, then $\boldsymbol{sel}(\mathbf{A}, w) \in \mathbf{A}$.
(centering) For all $\mathbf{A} \in P(\Omega)$ and all $w \in \Omega$, if $w \in \mathbf{A}$, $\boldsymbol{sel}(\mathbf{A},w) = w$.
Then: $[\![\text{will}_h \mathbf{P}]\!]^w = 1$ iff $[\![\mathbf{P}]\!]^{\boldsymbol{sel}(\mathbf{H}, w)} = 1$
[10] Example from Cariani (2021, 253).

2.3 The illocutionary solution

According to the illocutionary solution, the puzzle can be dissolved once we realize that, despite the appearance, Sam's utterances of (1) and (2) are not alike at all. (2) is a straightforward assertion, while (1) is a distinct speech act, like *prediction* (or *conjecture*). To predict that it will rain, according to this solution, is to say something like: it is likely that it will rain. Unlike assertion, prediction is subject to weaker norms. This explains why Sam's utterance of (1) is okay: he is merely making a prediction, so he does not need very strong evidence at all.

The illocutionary solution shares one of the problems for the modal solution: insofar as the solution locates the source of the puzzle at the specialness of future-directed statements, it has to commit to implausibly general evidential standards for all past-directed assertions. But as we have seen, not all past-directed assertions require causally-downstream evidence.

A different problem with the illocutionary solution is that it is unclear how we should understand prediction. If prediction is exactly like assertion except that it is subject to weaker norms, then the solution seems *ad hoc*; but if prediction is unlike assertion in other ways, it becomes hard to explain why many characteristic responses to assertion seem appropriate as responses to prediction as well. For example, we can typically challenge an assertion, but not a weaker speech act, with "How do you know?". Compare:

(11) It rained yesterday.
How do you know?
(12) I think it might have rained yesterday.
?How do you know?
(13) It will rain tomorrow.
How do you know?

In addition, the speaker who makes an assertion is typically required to retract their assertion, if what they have asserted turns out to be false.[11] If one is merely making a conjecture, however, there is no pressure to retract. Consider:

(14) Russ will come to the party.
(15) It is very likely that Russ will come to the party.

Intuitively, a speaker will need to retract (14) but not (15), if it turns out that Russ did not come to the party. These data suggest that, contrary to what the illocutionary solution says, a future-oriented assertion is still an assertion.

3 My solution: the pertinent evidence norm of assertion

What, then, is wrong with Sam's assertion of (1)? My proposal is that Sam is not in a position to assert (2) because he would be violating a norm of assertion, the *pertinent evidence norm*:

> (Pertinent evidence norm of assertion) One should assert that p only if one has pertinent evidence for whether p.

[11] See MacFarlane (2011).

Typically, we do not approach a line of inquiry like a blank slate: there are often distinctive *kinds* of evidence that we look for and expect for a given question in a particular context. For example, when scientists try to settle a scientific question, they design and conduct experiments; when philosophers try to determine whether a position is correct, they examine the arguments for and against that position; when detectives and judges try to deliver a verdict on a crime case, they look for witness testimonies and physical evidence. Following Daniel Friedman (forthcoming, p. 6), let's call such evidence "pertinent evidence." Pertinent evidence is evidence that we associate with certain kinds of inquiry, that we are inclined to look for and expect to find, and that we take to be *prima facie* necessary for settling the question at hand.

What makes some evidence pertinent? It seems that pertinent evidence is the *best normally available evidence*. In other words, for a kind of inquiry, a kind of evidence is pertinent in virtue of being the *best kind of evidence that we can reasonably expect an inquirer to acquire*, absent other information and considerations. Take a scientist's inquiry into whether an alloy is conductive. In this case, the pertinent evidence consists of results from scientific experiments. Some other kinds of evidence are best but not normally available: logical proof, or testimony from God; some kinds of evidence are normally available but not as good: a layperson's opinion, or the scientist's best guess based on their years of research experience.

It is important to note that in many cases pertinent evidence might in fact be unavailable or misleading. Perhaps the scientist is unable to conduct experiments at the moment, or the experiments would require technology too advanced for our time. Still, the fact that experimental results are in fact unavailable in a particular case does not change the fact that they are pertinent to the inquiry, since they are the kind of evidence that is normally available. The kinds of evidence that are pertinent to a kind of inquiry are relatively stable in the grand scheme of things.

Many take assertion to be subject to norms: i.e., there are conditions for appropriately asserting a proposition.[12] For example, KNA is widely accepted.[13] From the failure of the epistemic solutions, however, we have seen that KNA cannot help us with the expiring assertibility puzzle. The pertinent evidence norm, on the other hand, can. Since the kinds of evidence normally available for past- and future-directed inquiries differ, the kinds of evidence that are pertinent for them tend to differ as well. This explains why, in general, the same evidence can count as pertinent at an earlier time and not pertinent at a later time.

In particular, we might think that for inquiries into *Did it rain on Saturday?*, the weather report from Saturday is pertinent evidence. After all, the weather report is plausibly the best normally-available evidence. A direct observation might be even better, but we cannot reasonably expect that the inquirer just so happens to be in the right place at the right time; a weather forecast from Friday is normally available, but it is worse than the weather report from Saturday. So Sam's assertion of (2) calls for better evidence than the weather forecast. Pertinent evidence for inquiries about the future, on the other hand, is limited by the fact that certain kinds of evidence are *in principle* unavailable. For inquiries into *Will it rain on Saturday?*,

[12] Some (Williamson, 2000) endorse a stronger claim that assertion has *constitutive* norms: i.e., certain norms are not just true of assertion, but also part of what it is to make an assertion. In this paper, I do not endorse or discuss this stronger thesis.

[13] One exception is Lackey (2007), who argues that knowledge is not necessary for appropriate assertion.

pertinent evidence only includes things like the weather forecast, general reliability of weather forecasts, and so on. Since Sam does have evidence of these kinds, his assertion of (1) is okay.

My solution also explains why assertibility is stable in cases like **Vacation**. Typically, we do not expect friends to either report their itineraries to, or check in on, each other *all the time*. Of course, in reality we might be more or less close to our friends: we could be catching up once in a few years, or we could be talking everyday. But if we just pick a random person on the street, it would be unreasonable to expect them to have checked on some particular friend of theirs within the last few days. So while direct reports from Jack would be excellent evidence for deciding whether Jack spent the weekend in New York, it is not normally available and therefore not pertinent. Plausibly, the next best kind of evidence is sincere expressions of plans, which Alice does have. So Alice's utterance of (6) does not violate the pertinent evidence norm. More generally, we might think that expressions of plans and intentions are the best normally available evidence for deciding what particular individuals will do and have done.

In addition, my solution avoids a worry that Ninan (2022) anticipates for a different proposal, the *available-evidence solution*. According to this proposal, one's standing to assert depends not just on what evidence one has, but also on what evidence one *could get and does not have*. Without losing evidence, one could lose the standing to assert, because one can come to be in a situation where one *could* acquire new, better evidence. In **Rain**, therefore, Sam's utterance of (2) on Sunday is inappropriate because he could acquire stronger evidence by checking the weather report from Saturday, whereas such evidence is unavailable on Friday. Ninan is right to point out that this solution won't work, since the same puzzle arises even when we tweak the case such that the evidence available remains the same. Consider the following:

> **Destruction**[14]
> This case is exactly like Rain, except that right after Sam utters (1), an unknown cause wipes out the entire city (Sam happens to be out of town when all of this happens), destroying everything including any trace of possible rain, witness, or record. Sam learns about the destruction before Sunday, and for all everyone knows, the destruction has nothing to do with the weather.

Unlike in **Rain**, there is no new, stronger evidence to be acquired in **Destruction**. Nevertheless, the intuition remains that, whereas on Friday Sam is warranted to assert that it will rain, on Sunday he is not warranted to assert that it did rain. The available evidence solution falsely predicts that Sam's standing to assert would remain intact.

My view is immune to this worry, however. While both the available evidence solution and my solution concerns evidence, my solution concerns the kind of evidence that is *normally* available for a kind of inquiry, not the kind of evidence that is *in fact* available in a particular case. On my view, if a kind of evidence is pertinent for a kind of inquiry, it stays pertinent whether or not it happens to be available at the time of utterance. As a result, my solution makes the same prediction about **Destruction** that it makes about **Rain**: Sam loses the standing to assert that it rained, because he does not have evidence that is pertinent to the past-oriented inquiry of whether it rained.

So if the pertinent evidence norm is true, we have a solution to the expiring assertibility puzzle. But is there any reason to accept this norm, beyond its explanatory power? After all, we

[14] This is analogous to Ninan's **Death** case (Ninan, 2022, p. 20-1).

do not want to posit a new *ad hoc* norm of assertion anytime we encounter some problem cases. Fortunately, we are in a much better position than that. There is an independent explanation for the pertinent evidence norm of assertion. In the remainder of this section, I will argue that the pertinent evidence norm of assertion is a consequence of what an assertion is and how we should responsibly close inquiry. That is:

1. An assertion is a proposal to settle collective inquiry into a question.
2. One should not settle a collective inquiry unless one has pertinent evidence for that inquiry.
3. One should not propose something that one shouldn't do.
4. So, one should not make an assertion unless one has pertinent evidence for the relevant question.

3.1 Assertion as a proposal to settle collective inquiry

When we inquire on our own, we are engaged in *personal inquiries*: for example, when we think to ourselves, *Where did I leave my car keys?*. When we inquire collaboratively with other people, we are engaged in *collective inquiries*: for example, when a couple of friends try to decide which board game to play, when a team of detectives investigate a crime case, or when a group of scientists collaborate on a project to find out about the qualities of a new synthetic material. Of course, most of the time we engage in a mixture of both: personal inquiries rely on and inform collective inquiries, and vice versa.

Like personal inquiry, collective inquiry is a goal-directed activity: it is aimed at *settling a question for the group*. In personal inquiry, one aligns one's actions towards resolving a question for oneself; in collective inquiry, a group of people align their actions with the others' towards a shared goal of resolving a question for the group.

Now that we have a basic understanding of collective inquiry, there is a natural way to understand Stalnaker's (1974, 2002, 2014) influential notion of common ground: the common ground among a group of people is the set of collective inquiries that are closed for that group. I will begin by describing the two main theoretical roles that a notion of common ground is supposed to play: (i) explaining successful coordination, and (ii) tracking progress in conversations (and other forms of information exchange). I will show that this understanding of common ground is fit to play both of these roles.

Successful coordination requires a *shared and mutually recognized* body of information. Consider two cars at a four-way stop. Each of the drivers knows that the red car reached the stop sign before the blue car, but since neither of them knows whether the other knows that, they both hesitate to proceed. The information that the red car arrived first is shared but not mutually recognized to be so, and coordination fails as a result. Now suppose that the driver of the blue car waves. That allows the red car to proceed, since it now becomes mutually recognized that the red car can go first. Coordination, and collaborative activities in general, take place against a background of mutually recognized information.

This is one of the core insights that motivates Stalnaker's common ground framework. The common ground between a group of people is the body of information that is taken for granted, and mutually recognized to be taken for granted, among them. We can say that p is

common ground iff p is commonly taken for granted, i.e., everyone takes for granted that p, everyone takes for granted that everyone takes for granted that p, and so on.[15]

Intuitively, for a group of people to commonly take something for granted is for them to take collective inquiry about it to be closed or settled. If the drivers at the four-way stop take the question, *Will the red car go first?*, to be open, then they fail to take for granted that the red car will go first (or that the other is taking for granted that the red car will go first, and so on.) As a result, neither of them can proceed. Similarly, if the drivers commonly take for granted that the red car will go first, then there is no more collective inquiry to be done. It seems natural, then, to think of the common ground as the lines of inquiry that the group collectively takes to be closed.

The other role for common ground to play is to track the progress of our conversations. For Stalnaker, many of our conversations are collaborative activities aimed at collective inquiry. The common ground does not just make these conversations possible, but also gets updated by new conversational moves as the conversations evolve. Since these conversations aim at collective inquiry, common ground tracks progress in collective inquiry: i.e., which lines of collective inquiry have been resolved, and which are still open. Again, it seems natural and in keeping with Stalnaker's framework to think of the common ground as the set of collective inquiries that are closed.[16]

An important conversational move that one can make is assertion. Following Stalnaker (1978), I take asserting a proposition p to be a *proposal to make p part of the common ground*. This account of assertion understands it in terms of its characteristic effect.[17] When a speaker asserts that the meeting starts at 8, for example, if no one objects, we can expect it to become common ground that the meeting starts at 8. This allows the conversation to evolve in ways that presupposes that the meeting starts at 8.

If the common ground is a set of settled collective inquiries, then *an assertion is also a proposal to settle a collective inquiry*. Importantly, closing or settling a line of inquiry is different from merely suspending or abandoning it. Sometimes we have to *suspend* inquiry, perhaps in order to focus on other tasks: e.g., a group of friends might suspend inquiry into the question, *Why did the dinosaurs die?*, because they have to start inquiring into the more urgent question, *What should we make for dinner?*. Sometimes we choose to *abandon* a line of inquiry because we find ourselves unable to make progress: e.g., my parents and I might decide to abandon the question, *What did I receive for Christmas eleven years ago?*, because we think there is just no way to make progress. When we settle a line of inquiry, however, we take ourselves to have successfully completed inquiry, once and for all.

An assertion is a proposal to settle a collective inquiry, not just to suspend or abandon it. It purports to be the *final word* on a question. When one asserts that the meeting starts at 8, if

[15] The details of the iterated attitudes that constitute common ground are controversial. Proposals include common belief, common acceptance, common knowledge, and common knowledge about what is accepted. In this paper, I use "take for granted" as a neutral term for whichever attitude we find most suitable. For a detailed discussion, see Yalcin (2024).

[16] One way to formalize the idea that the common ground is a set of collective inquiries is to model the common ground as a set of question-answer pairs, as opposed to a set of possible worlds. The details of this proposal are beyond the scope of this paper. One thing to mention, however, is that this model makes common ground question-sensitive. This aligns with Yalcin's (2018) view that belief itself is question-sensitive.

[17] Other accounts of assertion understand assertion in terms of the attitude or epistemic position it expresses (Unger 1975), its constitutive norm (Williamson 2000), or the commitment one takes on in making an assertion (Brandom 1983). In principle, these accounts are not incompatible with the account I rely on here. For a detailed discussion, see MacFarlane (2011).

no one objects, we can expect the question, *When does the meeting start?*, to be settled for the group.[18] As a result, the conversation will evolve in ways that reflect the fact that this question is no longer open.

3.2 Pertinent evidence norm of inquiry

So assertion is a proposal to settle collective inquiry. When is it appropriate to settle an inquiry? Intuitively, whether one can appropriately settle a line of inquiry depends on one's evidence. While we would ideally like to have our evidence *entail* the answer, in practice it almost never does. Nevertheless, we treat many inquiries as closed: whether the earth orbits the sun, who won the last presidential election, why the sky appears blue, and so on. But of course, we cannot just close any inquiry we like, no matter how little evidence we have. The closing of inquiries, like other epistemic states and activities, is governed by norms, which determine when it is epistemically responsible.

In this section, I argue that the following is one of these norms:[19]

> (Pertinent evidence norm) One should not settle inquiry into a question unless one has pertinent evidence for that inquiry.

Note that this is a norm for *settling* inquiry, not for simply stopping inquiry. There is nothing epistemically irresponsible about stopping inquiry when one has little or even no evidence. The pertinent evidence norm does not say that one should keep actively inquiring if one lacks pertinent evidence. After all, there can be many good reasons not to inquire further: one might be out of time or resources, one might no longer find the question interesting, or one might have something more urgent to deal with. The pertinent evidence norm just says that one should keep the line of inquiry open.

The basic motivation for the pertinent evidence norm is a familiar idea: resolving some inquiries calls for certain *kinds* of evidence. For example, Buchak (2014) thinks that certain evidential bases can license a high credence without licensing a settled belief.[20] Consider the following pair of cases from Buchak:

> **Statistics**
> You leave the seminar room to get a drink, and you come back to find that your phone has been stolen. Barbara and Jake were the only ones in the room. You know that statistically men are 10 times more likely to steal than women. On this basis, you settled on the belief that Jake stole it.

> **Testimony**
> You leave the seminar room to get a drink, and you come back to find that your phone has been stolen. Barbara and Jake were the only ones in the room. You know that

[18] Typically, we can also expect a series of related questions to be settled as well. For example: *Does the meeting start at 8?*, *Does the meeting start before 9?*, and so on.

[19] Plausibly, another norm is the knowledge norm: An agent (or group of agents) ought not to close inquiry into a question unless they know the answer to that question.

[20] Buchak (2014) takes belief to be an all-or-nothing, question-settling attitude, to be distinguished from degrees of belief, or credence. There are reasons to think that belief is the wrong state to play this role. See Hawthorne, Rothschild, & Spectre (2016), Goodman (2023). Here I set aside issues about the right account of belief. For the sake of my argument, what matters is that Buchak's sense of belief is a question-settling state.

> statistically men and women are equally likely to steal, but a relatively reliable (90% reliable) witness tells you that they saw Jake steal it. On this basis, you settled on the belief that Jake stole it.

Intuitively, as Buchak points out, your settled belief in **Testimony** is appropriate whereas your settled belief in **Statistics** is not, even though your evidence makes the proposition you believe equally likely. The upshot is that the appropriateness of settling a question depends not just on the strength of one's evidence, but also on the kind of evidence.[21]

Which kinds of evidence are pertinent for an inquiry depend on the kind of inquiry it is, determined by a combination of rational considerations, convention, and a small amount of arbitrariness. As a result, we might be unable to articulate a precise formula for mapping each inquiry to its pertinent evidence. Nevertheless, mutual recognition of such pertinent evidence is an important part of collaborative activities. When we try to resolve a court case, for example, we expect things like witness testimonies, documents, and physical evidence, not just hearsay, character trait reports, or statistics. For many kinds of inquiry, which types of evidence are pertinent is typically common ground: experiments are pertinent for scientific inquiries, test results are pertinent for medical inquiries, and arguments are pertinent for philosophical inquiries, and so on. We take pertinent evidence to be *prima facie* necessary, though not always sufficient, for closing a line of inquiry, and to ignore pertinent evidence is epistemically irresponsible.

Why should there be such a special status for evidence, and why should we endorse the pertinent evidence norm of inquiry? Daniel Friedman (forthcoming) provides the following reason:[22] (i) one should not close inquiry unless one can reasonably expect that unpossessed evidence is low-impact, i.e., that unpossessed evidence won't require one to change one's mind, and (ii) one cannot reasonably expect pertinent evidence to be low-impact. After all, pertinent evidence is the best kind of evidence among those which are normally available, the kind that can easily outweigh other kinds of evidence.

The pertinent evidence norm is useful because complying with it allows an inquirer to better comply with the more nuanced reasons for and against closing inquiry but at a lower cost. We can compare two inquirers, Norma and Ada. Like most of us, Norma and Ada want to know the answers to many questions, but their resources for inquiry are limited. Norma and Ada are similar in how they inquire except for how to decide whether to close inquiry. Norma accepts the pertinent evidence norm: all else equal, Norma chooses to close inquiry only when she has the best normally available kind of evidence, and to leave inquiry open otherwise. Ada, on the other hand, does not rely on any norms. Each time she faces the decision of whether to close inquiry, she deliberates from scratch, weighing all the reasons for and against closing inquiry in that particular case.

---

[21] This is not the main upshot of Buchak (2014). Her main point has more to do with the connection between belief and reactive attitudes like blame and shame.

[22] To be more accurate, Daniel Friedman (forthcoming) endorses a stronger version of the pertinent evidence norm:

> (Strong pertinent evidence norm) If an agent (or group of agents) lacks evidence that is pertinent for an inquiry, they cannot appropriately close inquiry.

Daniel Friedman does not specify what makes some evidence pertinent. While I take us to be talking about the same notion, it is possible that he does not agree with my suggestion that pertinent evidence is the best normally available evidence.

Ada's case-by-case deliberation may, in principle, be more sensitive to fine-grained features of the situation. On the other hand, the advantage of Norma's strategy is that it allows her to approximate this sensitivity at a much lower cost. By deferring to the pertinent evidence norm, Norma avoids repeatedly re-evaluating whether her current evidence is sufficient given the stakes, the costs of further inquiry, the reliability of different evidential sources, and the risk of error. Instead, she relies on a norm that has been shaped to reliably track these considerations across a wide range of ordinary cases. Given the vast number of inquiries one needs to resolve and the scarcity of one's resources, we can expect Norma to be a more successful inquirer than Ada in the long run.[23]

3.3 Pertinent evidence norm of assertion

Typically, we should not propose something that would be inappropriate. So if my assertion would be a proposal to close collective inquiry on q, and it is inappropriate to close collective inquiry on q, then I should not make that assertion. And since it is inappropriate to close inquiry when we lack pertinent evidence, this results in a pertinent evidence norm of assertion.

This completes my argument for the pertinent evidence norm of assertion, as well as my solution to the expiring assertibility puzzle. But the implications of the pertinent evidence norm of assertion do not stop here. In particular, it seems to be able to explain two other observations.

First, some assertions strike us as odd, even when they are made by experts who in fact know what they assert. Consider:

> **Diagnosis**
> You are sitting in the office of Randy, your doctor. Last week you took some blood tests, but you don't know whether the results are ready yet. Randy walks in and tells you: "You have the condition." But unbeknownst to you, he did not learn that from the blood tests you took, since the test results are not ready yet. Instead, before you took those tests, Randy had already formed the belief that you have the condition on the basis of your symptoms and your demographics.

Something is off with Randy's assertion that you have the condition, but what exactly? One might think that Randy is not in a position to assert that you have the condition because he does not *know* that you do. Intuitively, however, it is not clear that Randy doesn't know. Experienced doctors come to know whether their patients have certain illnesses on the basis of their symptoms all the time. Why can't Randy come to know that you have the condition in the same way? In addition, it would not be so odd if Randy had said instead, "The test results are not ready yet, but I know that you have the condition, because …" It seems, then, that the oddness of Randy's assertion cannot be explained solely by whether he has knowledge.

Another possibility is that Randy's assertion is misleading: it suggests that he has learned about your condition from your test results, when he has not. But if knowledge is sufficient for assertion, why does it matter *how* Randy comes to know, given that he *does* know that you have the condition? And suppose Randy says instead: "The test results are not available yet, but you have the condition." This would still sound quite odd, even though it no longer communicates

[23] This story is different from but reminiscent of Ninan's (2022) explanation for why the standards for ruling out future possibilities are lower than those for ruling out past/present possibilities.

the misleading suggestion that Randy has seen the results. Why, then, does his assertion still sound strange?

Again, the pertinent evidence norm of assertion can help us out here. Randy does not yet have the test results, which are pertinent to medical inquiries. This is why his assertion that you have the condition is odd: in so doing, he violates the pertinent evidence norm of assertion.

The second observation that the pertinent evidence norm of assertion would explain is that we sometimes feel the need to make *evidentially hedged assertions*, as opposed to bare assertions. A hedged assertion is an assertion that has been *qualified*, and whose corresponding commitment has typically been weakened (Benton & van Elswyk, 2020).[24] One way to hedge an assertion is to indicate one's source of evidence, as in the following, for example:

(16) **Apparently**, Bill ate the cake.
(17) Bill **must** have eaten the cake.
(18) It is **possible** that Bill ate the cake.

Why do we sometimes choose to hedge assertions? Plausibly, it is because we are not in a position to make an assertion proper, because doing so would violate some norm of assertion. It is telling, then, that the hedges in (16)-(18) all indicate something about the *kind* of evidence that the speaker does *not* have, in addition to the kind of evidence the speaker does have: they all indicate that the speaker lacks direct perceptual evidence of the event in question. The pertinent evidence norm of assertion would offer us a natural explanation: direct visual evidence is pertinent for inquiries about particular, local events, and we hedge our assertions when we lack such evidence to avoid violating the pertinent evidence norm of assertion.

One might worry that these hedging data are better explained by KNA. If lack of certain kinds of evidence defeats knowledge, then one might choose to hedge to avoid violating the KNA, instead of the pertinent evidence norm. However, there are reasons to resist this idea. Mortini (2026) argues that attempts to justify the KNA from hedging overgeneralize. In addition, following von Fintel & Gillies (2010), I take "must" to be strong, i.e., must-p entails p. So if a rational agent knows that Bill must have eaten the cake, then they also know that Bill ate the cake. KNA cannot explain why, even when we know the conclusion of our inferences, we are still inclined to hedge our assertions with "must." More generally, we might think that testimonial evidence is a reliable source of knowledge. We come to know many things by trusting others' reports, and when we do so, we are inclined to evidentially hedge our assertions. We need the pertinent evidence norm of assertion to explain such cases.

3.4 My solution vs. existing solutions

We can now compare my solution to the three existing solutions. While they all face some problems, each of them gets something right about the expiring assertibility puzzle, and my solution preserves the core insights of each of them.

Ninan's (2022) and Torre's (2024) epistemic solutions offer an important insight about how we inquire: our standards can be lower when we inquire about the future and higher when we inquire about the past. In fact, Ninan has a story about why this practice might be

[24] It is controversial whether hedged assertions are always weaker than the corresponding bare assertions. Von Fintel & Gillies (2010) argue, for example, that "must-p" is just as strong as a bare assertion that p. In this paper, I follow the tradition of classifying must-statements as hedged assertions, but I remain neutral on whether hedged assertions always weaken the speaker's commitment.

evolutionarily advantageous.[25] However, they take this practice to affect knowledge, and only through knowledge, assertion. As a result, they have trouble explaining the linguistic data which suggest that the speaker does not lose knowledge in expiring assertibility cases. Unlike their approach, my solution recognizes assertion as itself constrained by our practices surrounding inquiry. By taking out knowledge as the middleman, we can preserve what is compelling about their observations while offering a more plausible explanation to the puzzle.

Cariani's (2021) modal solution correctly points out that assertions come with evidential constraints. But rather than making a sweeping claim that all non-modal assertions require causally downstream evidence, I suggest that different inquiries require different kinds of evidence, depending on which kinds of evidence are pertinent. I also agree with Cariani that future-directed assertions often require weaker evidence, but not because there is a special linguistic device that just cancels any evidential requirements. Rather, it is because of the nature of future-directed inquiries and, consequently, facts about the best normally available evidence for such inquiries. As a result, my solution is capable of explaining (9) and (10). Typically, we take many things to be common ground without ever inquiring into them. For example, we take it to be common ground that there is air, that the sun will rise again tomorrow, or that the earth did not shatter and then reform without anyone noticing. This means that we take collective inquiry into many questions to be settled without ever needing to inquire into them. It seems that both the question, *Did my brother just adopt a bunch of gorillas?*, and the question, *Did the sun rise from the east?*, belong to that category. So unlike many other assertions, (9) and (10) do not propose to close any collective inquiries that aren't already closed. In other words, asserting (9) or (10) does not propose any course of action that isn't already accepted by everyone else. As a result, the norms for closing collective inquiries are of no effect here. This explains why, contrary to what Cariani's (2021) proposal predicts, it is okay to assert (9) and (10) without any causally downstream evidence.

The illocutionary solution is right that the puzzle is puzzling partly because (1) and (2) appear to be the very same move, when they are not quite. However, it locates the difference in the kinds of speech acts they are. This leads to a dilemma: the more prediction resembles assertion, the less plausible and motivated the solution seems; but the less prediction resembles assertion, the harder it is to explain our linguistic behavior involving prediction. My solution, on the other hand, vindicates the intuitive idea that (1) and (2) are different moves in some sense, but not in the sense that they are not both assertions. Rather, they are different because they settle different lines of collective inquiries, even though they are assertions of the same proposition.

4 Conclusion

I have proposed a solution to the expiring assertibility puzzle by proposing and arguing for a new norm of assertion, the pertinent evidence norm. One upshot of this solution is that the norms governing inquiry are more multifaceted than we might have assumed. Whether we can appropriately close a line of inquiry depends not only on whether we know the answer to the question, but also on the kinds of evidence we have, and the kind of inquiry we are conducting.

[25] See Ninan (2022).

Another upshot of this solution is that it is a zetetic approach to a linguistic puzzle:[26] instead of thinking of assertion primarily as a way of transmitting information, we can situate it within the context of collective inquiry and understand it as a move with certain characteristic effects for updating and closing lines of collective inquiry. Consequently, the norms that govern assertion extend beyond the traditional epistemological concepts, like truth and knowledge, to a broader and richer set of zetetic norms.

More generally, I see my account as contributing to a larger shift towards placing inquiry at the center of theorizing about our cognitive and linguistic activities. As Jane Friedman (2020) argues, epistemologists should investigate questions about how to inquire, in addition to questions about what to believe. My solution to the expiring assertibility puzzle represents a similar strategy to philosophy of language: many conversational moves can be understood as inquiry-related moves. In this paper, I have focused on assertion, but it opens up the possibility of extending this approach to many other conversational moves or linguistic practices.

[26] By "zetetic" I mean related to inquiry.

References

Benton, M. & van Elswyk, P. (2020). Hedged Assertion. In Sanford Goldberg, *The Oxford Handbook of Assertion*. Oxford: Oxford University Press. pp. 245-263.

Brandom, R. (1983). Asserting. *Noûs* 17 (4):637-650.

Buchak, L. (2014). Belief, credence, and norms. *Philosophical Studies 169 (2)*, 1-27.

Cariani, F. & Santorio, P. (2018). Will done Better: Selection Semantics, Future Credence, and Indeterminacy. *Mind 127 (505)*, 129-165.

Cariani, F. (2021). *The Modal Future: A Theory of Future-Directed Thought and Talk*. Cambridge, UK: Cambridge University Press.

von Fintel, K. & Gillies, A.S. (2010) *Must* . . . stay . . . strong!. *Nat Lang Semantics* 18: 351–383.

Friedman, D. (2025). Intentions and Inquiry. *Mind*, *134 (533)*, 85-106.

Friedman, D. (forthcoming). Cooperation and Shared Inquiry. *Australasian Journal of Philosophy*.

Friedman, J. (2013). Question‐directed attitudes. *Philosophical Perspectives*, *27(1)*, 145-174.

Friedman, J. (2017). Why suspend judging? *Noûs*, *51(2)*, 302-326.

Friedman, J. (2020). The Epistemic and the Zetetic. *Philosophical Review 129 (4)*, 501-536.

Goodman, J. (2023). The myth of full belief. *Philosophical Perspectives 37 (1)*, 164-171.

Hawthorne, J.; Rothschild, D. & Spectre, L. (2016). Belief is weak. *Philosophical Studies* 173 (5):1393-1404.

Lackey, J. (2007). Norms of assertion. *Noûs* 41 (4):594–626.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74 (4), 549 – 567.

MacFarlane, J. (2011). What Is Assertion. In Jessica Brown & Herman Cappelen, *Assertion: New Philosophical Essays*. Oxford, GB: Oxford University Press.

Mortini, D. (2026). Evidentially Hedged Assertions and the Knowledge Norm of Assertion. *Erkenntnis* 91(6): 2781–96.

Ninan, D. (2022). Assertion, Evidence, and the Future. *Philosophical Review 131 (4)*, 405-451.

Ninan, D. (2025). Modality and the Future. *Analysis 85 (2)*, 515-527.

Stalnaker, R. (1974). Pragmatic presuppositions. In M. K. Munitz & P. Unger (Eds.), *Semantics and philosophy*. New York: New York University Press.

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics 9: Pragmatics* (pp. 315–332). New York: Academic Press.

Stalnaker, R. (2002). Common ground. *Linguistics and philosophy*, *25(5/6)*, 701-721.

Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.

Torre, S. (2024). Knowledge of the Future and Reliable Belief-Forming Processes. *Australasian Journal of Philosophy 103*, 39-57.

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

Williamson, T. (1996). Knowing and Asserting. *The Philosophical Review*, *105(4)*, 489–523.

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

Yalcin, S. (2018). Belief as Question‐Sensitive. *Philosophy and Phenomenological Research 97 (1)*, 23-47.

Yalcin, S. (2024). Defining common ground. *Linguistics and Philosophy 47 (6)*, 1045-1070.