Forthcoming in *Australasian Journal of Philosophy*

Knowledge of the Future and Reliable Belief-Forming Processes

Stephan Torre

Abstract

This paper embraces the view that we have substantial knowledge of the future and investigates how such knowledge fundamentally differs from knowledge of the past and present. I argue for a new source of context-sensitivity with respect to knowledge attributions arising from presuppositions about reliable belief-forming processes. This context sensitivity has important consequences for knowledge of the future, as well as the appropriateness of assertions about the future. I argue that not only is knowledge of future events typically brought about by fundamentally different processes from those that bring about knowledge of past events, that this is the case is often presupposed in attributing knowledge. I argue that this new source of context sensitivity naturally extends to explaining the recent puzzle of 'easy foreknowledge'.

Introduction

We know a lot about the future. Many of us know what we will do later today, that it will snow in Winnipeg next winter, that a dropped basketball will bounce. This paper examines such knowledge of contingent future events and identifies how it fundamentally differs from knowledge of contingent past and present events. I argue not only that knowledge of future events is typically brought about by fundamentally different processes from those that bring about knowledge of past events, but also that this fact is often presupposed in attributing knowledge. Presuppositions about reliable belief-forming processes introduce a new source of context sensitivity with respect to knowledge attributions.

In Section 1, I present and further develop the relevant alternatives framework introduced by David Lewis (1996). In 1.1, I extend Lewis's Rule of Reliability to include three processes by which we gain knowledge of future contingent events: induction, mental simulation, and practical foreknowledge. In 1.2, I argue that presuppositions about the kind of process that produced a subject's belief introduces a new source of context sensitivity with respect to knowledge attributions. In Section 2, I introduce Dilip Ninan's (2022) "Easy Foreknowledge" puzzle and his solution to it. I provide objections to Ninan's solution, and, in Section 3, I argue that the context-sensitive framework developed in Section 1 can accommodate cases of easy foreknowledge without making them essentially temporal. I then consider and respond to two recent objections from Fabrizio Cariani (2021) that have been leveled against epistemic accounts such as the one that I propose.

1 Knowledge, Relevant Alternatives, and Presuppositions about Belief-Forming Processes

I will argue for a new source of context-sensitivity for knowledge attributions by further developing David Lewis's relevant alternatives account as proposed in "Elusive Knowledge". Lewis states his account of knowledge as follows:

> **S** knows that P iff S's evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that we are properly ignoring (Lewis 1996, 554).

The knowledge attribution 'S knows that p' is true in a context just in case the subject's evidence eliminates all the not-p possibilities that are relevant (i.e. those not being properly ignored) in the context. Lewis stipulates 'evidence' to be one's "entire perceptual experience and memory" (Lewis 1996: 553) and I will follow Lewis in understanding the term in this way.[1] A possibility, w, is eliminated by one's evidence just in case the subject's perceptual experience

[1] Exactly how evidence is to be understood and how it relates to knowledge is a matter of substantial controversy. For example, Timothy Williamson famously equates a subject's knowledge with her evidence. Lewis's conception of evidence is clearly not understood in this way. I don't intend to take a stand on these issues by using 'evidence' in the way that Lewis does. I am happy to just stipulate that 'evidence', when left unqualified, is to be understood in Lewis's sense in this essay and leave it as an open question whether this use corresponds to the correct epistemological notion of evidence.

and memory in w fails to match her perceptual experience and memory in the actual world.

Lewis provides several rules of relevance for what possibilities can and cannot be properly ignored in a given context: the Rules of Actuality, Belief, Resemblance, Reliability, Method, Conservatism, and Attention. In further developing Lewis's framework to accommodate the kind of context-sensitivity I have in mind, I will focus on the Rule of Reliability and the Rule of Conservatism. In what follows, I will develop these rules of relevance in a particular way. The result is an account which is both a version of reliabilism and a version of the relevant alternatives approach. On such an account, what counts as the relevant reliable, belief-forming process depends on features of the context, and this in turn affects which possibilities are included and excluded in the set of relevant alternatives.[2]

1.1 Expanding the Rule of Reliability to include Non-Downstream processes

Lewis states the Rule of Reliability as follows:

> **The Rule of Reliability** "Consider processes whereby information is transmitted to us: perception, memory, and testimony. These processes are reliable. Within limits, we are entitled to take them for granted. We may properly presuppose that they work without a glitch in the case under consideration. Defeasibly—very defeasibly! — a possibility in which they fail may properly be ignored" (Lewis 1996: 558).

It is noteworthy that in discussing the Rule of Reliability, Lewis exclusively mentions causal processes: perception, memory, and testimony. Let us refer to these processes of belief formation collectively as *downstream processes* since they all involve a causal chain from the event that one has knowledge about to one's knowledge-constituting belief: there is a causal chain involving perception and memory from the event of it raining this morning to my knowledge-constituting belief that it rained. None of the reliabilist processes that Lewis cites in the

[2]There are certain comments that Lewis makes in stating the Rules of Reliability and Conservatism, which I cite below, that suggest an account along the lines that I develop here; however, I make no claim that the account developed here is what Lewis intended. Lewis's rules of relevance are loosely characterized and indeterminate, leaving them open to further precification and development.

Rule of Reliability can apply to knowledge of future events for the obvious reason that subjects are not causally downstream of such events. Reflecting on the nature of our knowledge of future events, it seems that even though our knowledge-constituting beliefs are not causally downstream of the events that our beliefs are about, there are several reliable processes by which we acquire such beliefs. Below I propose expanding the Rule of Reliability to include three such reliable, non-downstream processes of belief formation: induction, mental simulation, and practical foreknowledge.[3]

In the case of inductive reasoning, we reason from a finite number of observed instances of event-type A having property F, to unobserved instances of events of type A having F. The unobserved instances might be in our past, such as when we reason inductively that it snowed in Winnipeg in 1872 after observing that it snowed every winter in Winnipeg since records began in 1873. Or the unobserved instances might be in our future, such as when we reason that it will snow in Winnipeg next winter. Interestingly, in the case of inductive reasoning about the past, there is oftentimes a competing downstream process for forming the belief about the past event. We might form the belief that it snowed in Winnipeg in 1872 inductively since it has snowed in Winnipeg every other year since records began. Or we might form the belief via a downstream process: perhaps we find historical accounts of the winter of 1872 in Winnipeg which mention that it snowed. In the case of beliefs about future events such as that it will snow in Winnipeg next winter, there is no competing downstream process for belief formation; only non-downstream processes are available.

Many of our beliefs about the future are based on a particular kind of inductive reasoning: causal inference. Since causes precede their effects, knowledge of a past or present cause can provide knowledge of a future effect. In defending the causal theory of knowledge, Alvin Goldman (1967) countenances foreknowledge based on inference from a common cause. For example, one might know that it will rain soon by observing the gathering rain clouds moving in from the west. The event of it raining soon is not a cause of the belief that it will rain. However that it will rain soon and the subject's belief have a common cause in the gathering rain clouds moving in from the west. In such a case the subject reasons inductively from the common cause to a future effect.

Another way in which we form beliefs, and gain knowledge, about future

[3]It is worth noting that Lewis includes induction under a different rule, a Rule of Method, and acknowledges that this rule may be subsumed under the Rule of Reliability. See Lewis (1996: 558-59).

events is by mental simulation. From across the pub, you see an empty glass get knocked over, and roll towards the end of the bar. You form the beliefs that the glass will continue its trajectory, fall off the end of the bar, and shatter. Your beliefs are true and, given the right conditions, may constitute knowledge. It is possible that these beliefs about contingent future events are derived from inductive reasoning: you've witnessed numerous cylindrical objects roll and fragile objects shatter when dropped. And you might use this inductive base to reason for the conclusion that the glass will roll to the end of the bar, fall, and then shatter. But there are compelling reasons to be doubtful that this is the whole story of how you gain such knowledge. An increasingly popular account from contemporary cognitive science is that mental simulation plays a substantial role in the formation of our beliefs about the future. According to this account, we possess an "intuitive physics engine", a cognitive model of how physical objects in our environment behave, which allows us to predict the future state of objects in our environment with substantial accuracy.[4]

In addition to simulating how objects will evolve based on an intuitive physics engine, we also have something like an intuitive psychology engine. We simulate the mental states of others and thereby their future mental states and actions. You need to cancel plans with a friend tomorrow and you simulate her learning of your cancellation. You conclude correctly that she will be disappointed. A plausible account claims that you gain knowledge of her mental state by putting yourself in her shoes, using your imagination to view the event from her point of view. As many have argued (Goldman 2006, Nichols and Stich 2003), mental simulation of this sort plays a substantial role in our knowledge of others' mental states and actions and it does not seem that inductive inference can provide the whole story. This process of acquiring knowledge also extends to foreknowledge. Assuming the right epistemic features obtain (for example, your mental simulations are sufficiently reliable), beliefs formed in such a way

[4]There is a substantial recent literature in cognitive science on mental simulation. Some good representative examples are Schwartz and Black (1999); Battaglia, Hamrick and Tenenbaum (2013); and Hamrick (2019).There is an interesting and significant epistemological question about whether mental simulation by imagination provides a new source of knowledge that is not reducible to observational and inferential knowledge. Sarah Aronowitz and Tania Lombrozo (2020) convincingly argue that "mental simulation cannot be reduced to observation or inference, though it shares important similarities with both" (1). In light of the recent literature on the topic, I'm inclined to agree with Aronowitz and Lombrozo that simulation is a distinct source of knowledge, irreducible to observation and inductive inference, and so it is a non-downstream belief-forming process distinct from inductive inference. For a more in depth discussion of mental simulation as a means of providing foreknowledge, and associated debates, see Cariani (2021, Chapter 12)

can and do constitute knowledge of the future.

A third way by which we gain knowledge of the future, which does not appear to be analyzable in terms of induction or mental simulation, is what Richard Moran (2004) calls "practical foreknowledge". I know many things about what will happen later today largely because I intend to make them happen. I intend to go on a run, I intend to take my daughter to the playground, and I intend to reply to emails. The way by which I come to have this knowledge of future events does not seem to involve induction or mental simulation. Suppose my partner comes to believe that I will go running this afternoon on the basis of observing that I run almost every sunny afternoon and this afternoon is sunny. My partner's knowledge seems fundamentally different from mine in that my decision and intention to go running is the basis for my knowledge, rather than the inductive support (I could plausibly know that I will go running even if I am unaware of the inductive support on which my partner's belief is based).[5] Assuming the right epistemic features obtain (one succeeds in doing what one intends to do, one reliably does what one intends to do, etc) corresponding beliefs about what one intends to do can constitute knowledge of the future.

I propose expanding Lewis's Rule of Reliability to include these reliable, yet non-downstream, processes for belief-formation. Assuming induction, mental simulation, and forming intentions can count as reliable belief forming processes, they are to be included within the Rule of Reliability. Within limits, defeasibly, and subject to other contextual features, we are entitled to take these reliable belief-forming processes for granted.

1.2 Presuppositions about Reliable Belief-Forming Processes

Let us now turn to the Rule of Conservatism. Lewis states the rule as follows:

The Rule of Conservatism "Suppose that those around us normally do ignore certain possibilities, and it is common knowledge that they do, (They do, they expect each other to, they expect each other to expect each other to, ...) Then—again, very defeasibly!—these

[5]The case of my knowledge that I will go running and my partner's knowledge of the same event is an example of the distinction Anscombe (2000: 56) highlights between intention and prediction. She famously considers the case of a man with a shopping list and a detective observing the shopper. They may both have knowledge of what the man will purchase, but one's knowledge is based on his intention to purchase certain items and the other's is based on predictions of behavior.

> generally ignored possibilities may properly be ignored. We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us" (Lewis 1996: 559).

Lewis goes on to note that the Rule of Conservatism can be understood in such a way that it subsumes the Rule of Reliability. He states, "we might subsume the Rule of Reliability under the Rule of Conservatism, on the ground that the reliable processes whereby we gain knowledge are familiar, are generally relied upon, and so are generally presupposed to be normally reliable" (559). If we follow Lewis's suggestion and include the Rule of Reliability under the Rule of Conservatism and understand the Rule of Reliability sufficiently broadly to include reliable processes generally, a new source of context-sensitivity with respect to knowledge attributions emerges. In introducing the Rule of Conservatism, Lewis writes "Suppose that *those around us* normally do ignore certain possibilities, and it is *common knowledge* that they do, ... We are permitted, defeasibly, to adopt the usual and *mutually expected presuppositions of those around us*" (Lewis 1996: 559, my emphasis). If we understand "those around us" as changing across contexts and we allow that what serves as common knowledge can also vary across contexts, then the Rule of Conservatism introduces a source of context-sensitivity for knowledge attributions: a certain process may be presupposed to be the reliable process by which the subject gains information in one context and not presupposed to be a sufficiently reliable process in another context, leading to variability with respect to knowledge attributions.

This context variability with respect to knowledge attributions is important for understanding knowledge of the future vs knowledge of the past. The current proposal involves acknowledging certain effects that conversational presuppositions can have on the set of alternatives relevant for assessing knowledge claims.[6] In Lewis's statement of the Rule of Conservatism he seems to explicitly endorse the view that what is "common knowledge" among the conversational participants plays a role in determining what possibilities can be properly ignored and provides a gloss of the iterative conception of common ground. This is a notion that has been developed in detail by Robert Stalnaker (1974, 1978, 1999, 2002). The conversational common ground is a set of propositions that are

[6] Another development of Lewis's relevant alternatives account in which conversational presuppositions play a role in determining the set of relevant alternatives is provided in Michael Blome-Tillman's *Knowledge and Presuppositions* (2014). Blome-Tillman proposes replacing Lewis's Rule of Attention with a presuppositional effect according to which those alternatives that are conversationally presupposed, rather than salient, cannot be properly ignored. He argues that such an account is better equiped to respond to skeptical challenges.

presupposed by the conversational participants. On the Stalnakerian account, a proposition is presupposed by the speaker just in case the speaker accepts the proposition, he accepts that his conversational participants accept it, he accepts that his conversational participants accept that each member of the conversation accepts it, etc. For example, suppose Susan and Tom are standing outside in the sun conversing. In their conversation, it may be presupposed that it is sunny out: Susan accepts that it is sunny, she accepts that Tom accepts that it is, she accepts that Tom accepts that Susan accepts that it is sunny, and so on.

The current proposal is that, in a given conversational context, there are presuppositions about what kind of reliable belief-forming processes are in play for a given proposition. Consider the following two contexts:

PAST: On Tuesday, Susan is talking to Tom, and she says, "It rained this morning".

FUTURE: On Wednesday, Susan is talking to Tom, and she says, "It will rain this afternoon".

Let us assume that in both contexts the proposition that Susan asserts is true. We can note some interesting and relevant differences in what is presupposed in both contexts. In the first context, it seems plausible that it is part of the conversational common ground that the belief that Susan asserts is formed by downstream processes: some combination of perception, testimony, and memory. In the absence of any qualifications or hedges, it would seem inappropriate for Susan to assert "It rained this morning" if her belief that it rained was formed by induction and she had no downstream evidence from the event of it raining this morning. If her belief was acquired in a non-downstream way: say, by checking the weather yesterday and inductively inferring that the weather forecast was correct, she ought to hedge in some way: "It probably rained this morning. The forecast is usually correct". The fact that an assertion of "It rained this morning" would be inappropriate in the absence of downstream evidence suggests that appropriately asserting that it rained earlier this morning typically presupposes that the belief was formed by a downstream process.

However, in FUTURE it is common knowledge among the conversational participants that the belief that Susan asserts is *not* formed entirely by downstream processes. Rather it is presupposed that such beliefs are formed through inductive evidence: perhaps by looking at the weather forecast, knowing that the forecast is reliable, or by inferring inductively from a particular cloud formation, for example. It is common knowledge that reliable belief-forming processes

about future events are non-downstream processes; whether implicitly or explicitly, we all know that beliefs about future events are not causally downstream of the events that the beliefs are about.

Let us return to the Principle of Conservatism: "We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us" (Lewis 1996: 559). The presuppositions at play in PAST differ from the presuppositions at play in FUTURE, and so the possibilities that are ignored, and permitted to be ignored, differ in the two cases. In PAST, it is presupposed that Susan's belief is formed by downstream processes and so possibilities compatible with her evidence in which either perception, memory, or testimony is faulty, and it did not rain in the morning can be properly ignored. In FUTURE, however, it is not presupposed that her belief is formed by a downstream process; instead, it is presupposed that her belief is formed by some non-downstream process, likely induction. None of the relevant alternatives are ones in which Susan's belief is formed exclusively by perception, memory, or testimony. Rather the possibilities relevant for evaluating Susan's knowledge are ones in which her belief is formed in part by a non-downstream process such as induction and we are permitted to ignore worlds in which induction is faulty and it does not rain this afternoon. The difference in what is presupposed about the reliable process for forming the belief asserted in PAST vs the reliable process for forming the belief asserted in FUTURE leads to a difference in which possibilities can be properly ignored and hence which possibilities are relevant for whether Susan knows in a particular context the proposition that she asserts.

In the case of beliefs about future events, there is a presupposition that the belief was not formed entirely by downstream processes. In the case of beliefs about past events, there is frequently—but not always!—a presupposition that the belief was formed by downstream processes. Whether such a presupposition holds depends on the event in question and features of the context. Consider the following case:

> **Cecilia's Departure**
>
> Aidan and Bianca have spent the afternoon together preparing for the arrival of their friend Cecilia who is travelling by train from New York. They both know that the train was scheduled to depart at 1pm. Suppose that Bianca looks at her watch, sees that it is 2pm, and says to Aidan, "Cecilia has departed already".

The belief that Bianca asserts is true, and the context is one where Aidan knows

(and Bianca knows that Aidan knows, and Aidan knows that Bianca knows that he knows, etc...) that Bianca's belief that Cecilia has departed is not formed entirely by downstream processes. Given that they have been together for the afternoon, Aidan knows that Bianca hasn't heard from Cecilia that she has departed, nor has she seen her depart. Aidan knows (and Bianca knows that he knows) that Bianca's belief formation involves, at least in part, non-downstream processes such as induction (the train Cecilia is travelling with typically departs on time) and perhaps also mental simulation (Cecilia said she will take the 1pm train from New York, Cecilia is trustworthy and reliable; she typically does what she says she will do). And so in ascribing knowledge to Bianca in this context, not-p possibilities in which the train is uncharacteristically delayed or Cecilia exhibits atypical behaviour may be properly ignored. Given the context, Bianca may count as knowing a past event (that Cecilia departed) even though the belief was not acquired entirely by downstream processes.

Suppose that another friend, Dieter, calls Bianca and asks whether Cecilia has departed. Here it may well be inappropriate and potentially misleading for Bianca to respond 'Yes, she has'. It would be better for Bianca to hedge and say something like, 'Probably; her train was scheduled to leave an hour ago', thereby indicating that she has no direct downstream evidence of her departure. Dieter, unlike Aidan, does not share with Bianca the presupposition that her belief was formed by non-downstream processes. This fact explains why it would be potentially misleading for Bianca to respond affirmatively, and why it would be better to hedge. In such a context, possibilities in which the same inductive and mental simulation processes are used, yet the train departed late, or Cecilia missed the train are potentially relevant. Whether Bianca counts as knowing in a given context is determined to some extent by the presuppositions about which belief forming processes are in play..[7]

Another example suggestive of the same kind of context sensitivity:

Out of Milk

Enzo and Felix are driving home together when Enzo says, "Oh no!

[7] An alternative diagnosis of the case is that Bianca knows that Cecilia has departed in both cases, however the hedging when speaking to Dieter is used to signal the absence of downstream evidence for the claim. However, such a diagnosis fails to explain why an unhedged assertion of "Yes, she has" seems inappropriate despite maintaining that Bianca knows the relevant proposition. The proposal here that she fails to know the relevant proposition in this context has the advantage of being able to explain the inappropriateness in terms of the knowledge norm of assertion: it would be inappropriate because she fails to know the relevant proposition. Thanks to an anonymous referee for suggesting this alternative diagnosis.

We're out of milk!".

"No worries", replies Felix, "The supermarket up ahead has milk".

Let us suppose that the belief that Felix asserts is true, and it is formed by induction: he has been to this supermarket several times in the past, successfully purchased milk there, and inductively infers that they now have milk. Furthermore, the context is such that it is common knowledge that Felix does not have direct perceptual or testimonial evidence that the supermarket currently has milk; they are in the car, Felix hasn't communicated with anyone from the store, etc. It is presupposed in this context that Felix's belief is formed by induction. Within this context, Felix's inductive reasoning can be taken for granted. Worlds in which his belief is formed by this method and the supermarket does not currently have milk may be properly ignored. In such a context, it seems plausible to claim that Felix knows that the supermarket up ahead has milk. There a plenty of ordinary situations relevantly like this one in which we take ourselves and others to know the relevant proposition.

After arriving at the supermarket, Felix goes inside to buy milk while Enzo waits in the car. After a few minutes, Enzo phones Felix and asks "Does the supermarket have milk?" If Felix hasn't yet reached the milk aisle, it would be odd for him to reply, "Yes, they do". It would be more appropriate for him to reply something like, "Hold on, I've only reached the bread aisle. Give me a minute!" The context has changed in such a way that the relevant belief-forming process is a downstream process: likely perceptual (seeing the milk) or perhaps testimonial (asking a store clerk if they have milk). In such a context, worlds in which Felix has the same inductive evidence but the store does not have milk are no longer properly ignored.

To sum up, what belief forming processes are presupposed in a given context introduces variability with respect to whether a subject counts as having knowledge of the event in question. Whereas it is presupposed that knowledge of future events always involves beliefs formed by non-downstream processes, in the case of knowledge of past events, while typically it is presupposed that such beliefs are formed by downstream processes, there are contexts in which it is presupposed that the belief is formed by non-downstream processes. Presuppositions about which belief-forming processes are in play in a particular context can affect whether a subject counts as knowing the relevant proposition in that context.

2 Ninan's Easy Foreknowledge Puzzle

In "Assertion, Evidence and the Future" Dilip Ninan introduces a fascinating new kind of puzzle involving assertions and knowledge about the future. Ninan argues for a new kind of case in which a subject loses knowledge over time without gaining or losing evidence. Below is Ninan's presentation of one such case:[8]

> **The Beth Case**
>
> Andy is a personal chef to a wealthy entrepreneur, Beth. Andy is making a new dish for Beth's dinner tonight (suppose it is a Friday). Based on his knowledge of the sorts of foods that Beth usually likes, Andy says to his friend Chris,
>
> (1) Beth will like this when she eats it.
>
> Andy finishes preparing the dish, and heads home for the night, before Beth gets back from work to eat dinner. When Beth returns, she eats the dish Andy has prepared, and thoroughly enjoys it. The next morning (Saturday), one of Andy's friends asks Andy, Did Beth enjoy the dish you made for her yesterday? Andy hasn't heard from Beth or anyone else whether or not she enjoyed the dish. I think it would seem odd here for Andy to flat-out assert that Beth liked the dish, i.e. to say,
>
> (2) Yes, she liked it.
>
> In order to make that claim, Andy would need to be more directly connected to the fact that Beth enjoyed the dish in question. For example, Andy would need to have been told by Beth or someone else that she did in fact enjoy the dish. Absent evidence of that sort, it would be better for Andy to hedge in some way, i.e. to say one of the following:
>
> (3) She probably enjoyed it.
>
> (4) She must have enjoyed it—it was just the sort of thing she usually likes. (Ninan 2022: 1-2)

Ninan goes on to argue that on Friday Andy knows that Beth will enjoy the dish, however on Saturday morning he does not know that Beth enjoyed the dish.

[8] Ninan presents a few such cases, however this example is the one that his discussion focuses on, and I will follow suit.

This difference in knowledge between Friday and Saturday morning explains the appropriateness of (1) and the inappropriateness of (2). In such cases, termed "easy foreknowledge" cases by Ninan, it is easier to obtain knowledge of a proposition that concerns the future than it is to obtain knowledge of the same proposition when it concerns the past.

Ninan's argument for easy foreknowledge from cases such as the Beth Case involves some well-defended intermediate conclusions. For ease of exposition, I will refer to the time of Andy's Friday assertion as 't_{BEFORE}' and the time the following morning when Andy is asked whether Beth enjoyed the dish as 't_{AFTER}'. I will also follow Ninan in taking the proposition that Andy expresses at t_{BEFORE} with (1) to be the very same proposition that Andy would affirm at t_{AFTER} with (2). I will refer to this proposition as 'ENJOY'.[9] Ninan endorses the Knowledge Norm of Assertion: that asserting p is appropriate only if one knows that p.[10] He proposes that the best explanation of what makes (1) appropriate and (2) not appropriate is the fact that Andy knows ENJOY at t_{BEFORE} and does not know ENJOY at t_{AFTER}.

Largely by being persuaded by his arguments, I agree with Ninan that Andy knows ENJOY at t_{BEFORE} and Andy does not know ENJOY at t_{AFTER}, and this difference in knowledge explains the difference in appropriateness in assertibility between (1) and (2). I expect, however, that many will seek an explanation elsewhere.[11] However, Ninan provides several compelling reasons in favor of the claim that Andy knows ENJOY at t_{BEFORE} and fails to know ENJOY at t_{AFTER}. Not only does it provide an explanation of the assertibility of (1) and the unassertibility of (2), it also accords well with other norms associated with knowledge such as those involving practical reasoning and interrogative attitudes (See Ninan 2022, 10-11). For the remainder of the paper, I will grant Ninan's claim that Andy knows ENJOY at t_{BEFORE} and fails to know it at t_{AFTER}, and

[9] An alternative view maintains that Andy expresses a tensed proposition at t_{BEFORE} with (1) and would affirm a distinct tensed proposition at t_{AFTER} with (2). That the same proposition is expressed with (1) and affirmed with (2) is controversial, but, I think correct. See Ninan (2022, Section 2) for a defense of this claim.

[10] The Knowledge Norm of Assertion is argued for in Williamson (2000) and Hawthorne (2004) among others.

[11] An alternative line of response is to deny that there is a difference with respect to Andy's knowledge of ENJOY at t_{BEFORE} vs t_{AFTER}: one might either claim that he knows ENJOY in both contexts, or he fails to know ENJOY in both contexts. One might then go on to give an alternative story about why an assertion of (1) is appropriate and an assertion of (2) is inappropriate. One way of doing so might be by appealing to some pragmatic differences between assertions of (1) and assertions of (2). Or differences in lexical meaning between 'will' claims like (1) and past claims like (2). A response of this latter sort is defended by Fabrizio Cariani in Chapter 13 of *The Modal Future*.

I will show that this phenomenon fits naturally into the contextualist framework outlined in the previous section.

2.1 Ninan's Preferred Solution to the Puzzle: Future Normality

Ninan ultimately concludes that "time matters" in explaining the Beth Case and similar ones. His proposed solution involves an appeal to future normality: in short, we "enjoy a default entitlement to assume that the future will unfold in a relatively normal manner" (Ninan 2022: 17). Ninan also uses as his starting point the relevant alternatives account proposed by Lewis in "Elusive Knowledge". To provide an explanation of "easy foreknowledge" cases like the Beth Case, Ninan proposes to add the following rule of relevance to Lewis's rules:

Rule of Future Normality Other things being equal, a possibility v is relevant for an agent x at time t in world w just in case: (i) v is *prima facie* relevant for x at t in w, and (ii) v's future unfolds in a way that is at least as normal as w (Ninan 2022, 24).

This new rule requires some explanation. To claim that a world is "*prima facie* relevant" for x at t in w, is to claim that it is relevant according to the rules of relevance that Lewis proposes in "Elusive Knowledge". In other words, if a world is among the set of relevant alternatives in a given context according to the Rules of Actuality, Belief, Resemblance, Reliability, Method, Conservatism, and Attention, then it is *prima facie* relevant.

Ninan's new rule appeals to the recently influential notion of normality. The notion of a future unfolding in a normal manner can be understood in a largely intuitive manner. A future in which it snows in Winnipeg on Christmas would be normal, however one in which it is a balmy, 35 degrees C in Winnipeg on Christmas would be very abnormal. A future in which I go on a run this afternoon would be largely normal, whereas a future in which I rob a bank this afternoon would be very abnormal. Martin Smith, who appeals to normality in his account of justification, links it with explanation: "normal conditions require *less explanation* than abnormal conditions do" (Original emphasis, Smith 2010: 15). Abnormal situations "cry[] out for *explanation* of some kind" whereas normal situations do not (Original emphasis, Smith 2010: 15). Whereas snowing on Christmas in Winnipeg wouldn't require much in the way of explanation, being a balmy 35 degrees C would presumably require substantial explanation.

The second condition in Ninan's Rule of Future Normality makes use of a two-place relation between worlds: that of x unfolding in a way that is at least as normal as y. Whereas the world in which I go on a run this afternoon is at least as normal as the world in which I rob a bank, the world in which I rob a bank is not at least as normal as the world in which I go on a run. This means that what worlds unfold in a sufficiently normal manner is relative to how things unfold in the actual world.

How does the Rule of Future Normality accommodate cases of easy foreknowledge exhibited in cases like the Beth Case? At t_{BEFORE}, all the worlds relevant for evaluating whether Andy knows ENJOY are worlds compatible with Andy's evidence at t_{BEFORE} and evolve normally after t_{BEFORE}. The Rule of Future Normality excises *prima facie* relevant worlds with abnormally evolving futures from the set of relevant alternatives. ENJOY is true in all remaining worlds, so Andy knows ENJOY at t_{BEFORE}.

What about at t_{AFTER}? The Rule of Future Normality only excises worlds that evolve abnormally *after* the time of evaluation. So worlds with abnormally evolving pasts relative to t_{AFTER}, such as worlds in which the dish spoils, or Beth gets ill, may well be among the alternatives relevant for assessing Andy's knowledge. These are worlds compatible with Andy's evidence in which ENJOY is false. So, the existence of such worlds in the set of relevant alternatives entails that Andy fails to know ENJOY at t_{AFTER}.

So, all else being equal, knowledge of the future may be easier to come by than knowledge of the past because the set of alternatives relevant for assessing knowledge of the future is restricted to those with normally evolving futures. When assessing a subject's knowledge of the past, no such restriction applies, and so abnormal pasts may be relevant and may well undermine knowledge.

2.2 Problems for the Future Normality Account

Ninan's explanation for why Andy knows ENJOY at t_{BEFORE} and fails to know ENJOY at t_{AFTER} is because at t_{AFTER} abnormal possibilities in which not-ENJOY are included among Andy's relevant alternatives that were properly ignored at t_{BEFORE}. Such not-ENJOY possibilities, which are uneliminated by Andy's evidence, are ones in which the dish spoiled, or Beth became ill and her illness caused her not to enjoy the dish. But if such possibilities are sufficient for explaining why Andy fails to know ENJOY at t_{AFTER}, it seems that there are similarly nearby and relevant possibilities at t_{BEFORE} that evolve normally

after t_{BEFORE} in which ENJOY is false. Call such worlds 'Normally Evolving Abnormal Past' or NEAP worlds. A world is a NEAP world at t, just in case something abnormal happens prior to t and it evolves normally after t. In one such NEAP world, Beth contracts a virus on Thursday which normally evolves through to her arrival home from work on Friday night when she does not enjoy the dish due to her illness. In another NEAP world, some *e coli* gets into the ingredients on Thursday which then normally evolves to spoil the dish at the time when Beth eats it. Ninan glosses the comparison of relative normality in how two worlds unfold as follows "the way [v] develops after t, given what happened in [v] up until and including t, might be *more normal than* or *less normal than* or *as normal as* the way w develops after t, given what happened" (Original emphasis, Ninan 2022, 23). So given what happens up to and including t_{BEFORE} in the world in which Beth contracts a virus on Thursday, a normal way for the future to unfold would be one in which she becomes ill and does not enjoy the dish.[12] If the explanation for why Andy fails to know ENJOY on Saturday morning is that abnormal past possibilities in which ENJOY is false become relevant, then possibilities with abnormal pasts at t_{BEFORE} that evolve normally after t_{BEFORE} ought to be relevant at t_{BEFORE}. I see no non-*ad hoc* way of allowing abnormal past alternatives in explaining why Andy fails to know at t_{AFTER} and not allowing NEAP possibilities at t_{BEFORE} (Similarly, considerations for leaving out NEAP worlds at t_{BEFORE} would also serve as considerations for leaving out abnormal past worlds at t_{AFTER}). However, if we allow NEAP possibilities into the set of relevant alternatives at t_{BEFORE}, we get the result that Andy fails to know ENJOY at t_{BEFORE} and he fails to know ENJOY at t_{AFTER}. The account fails to accommodate the easy foreknowledge exhibited by the Beth Case.

Another difficulty with the future normality account comes from cases like the ones presented in ï¿œ1.2: Cecilia's Departure and Out of Milk. They seem to exhibit a similar phenomenon to the Beth Case, however they do not involve knowledge of the future, instead they involve knowledge of the past (that Cecilia departed) or the present (that the supermarket has milk). Therefore, the Rule of Future Normality cannot be applied to explain these cases. In fact, given that the event in question is not in the future, Ninan's account seems committed to

[12] It is important to note that we are not comparing whether the future of v is as normal as the future of w *simpliciter*, but rather whether the future of v given what has happened in v up to and including t is as normal as the future of w given what has happened in w up to and including t.

maintaining that there is no change in knowledge from the earlier to the later context, and hence no change in assertability.[13] One option for the proponent of the future normality account is to maintain that the Rule of Future Normality explains the Beth Case, however Cecilia's Departure and Out of Milk exhibit a different phenomenon with a different explanation. But such a response seems implausible. The phenomena exhibited by Cecilia's Departure and Out of Milk *feel* the same as that exhibited by the Beth Case.

I think these cases involving past or present knowledge show that the phenomena exhibited in these cases is not essentially temporal. Rather it comes from a context sensitivity concerning shared presuppositions about reliable belief-forming processes. As a result, the phenomenon of easy foreknowledge is *derivatively temporal*: there is a presupposition that beliefs about the future are formed by non-downstream processes, and typically (but not always!) no such presupposition exists in the case of beliefs about the present or the past.

This leads to a more general complaint with the future normality response. There is a sense in which it seems explanatorily deficient. A natural question to ask is: "Why do we remove abnormally evolving futures when evaluating knowledge of the future and not remove abnormally evolving pasts when evaluating knowledge of the past?" What temporal feature underlies this asymmetry? Positing it as a brute asymmetry in our knowledge claims seems unsatisfactory. At least, an account that was able to provide an explanation for easy foreknowledge cases that didn't posit this brute asymmetry would be preferable.

3 An Alternative Solution: Presuppositions about Reliable Belief-Forming Processes

Rather than adding a new rule of relevance, the Rule of Future Normality, which introduces a temporal asymmetry into our knowledge of past and future events by fiat, the Beth Case and similar examples can be explained by the more general kind of context sensitivity outlined in ï¿œ1. Easy foreknowledge cases are not essential temporal, rather they fall out of more general considerations concerning presuppositions about what kinds of processes are involved in belief formation. The fact that the belief-forming processes that provide knowledge of future events are typically of a different kind than the belief-forming processes that provide knowledge of past events, and that this is common knowledge, explains

[13] Cf. Ninan's discussion of "stable foreknowledge" in ï¿œ6.5.

the Beth Case and similar cases while positioning them within a broader theory and doesn't make the feature essentially temporal.

Andy knows ENJOY at t_{BEFORE} because at t_{BEFORE} it is presupposed that the process by which Andy formed the belief that ENJOY is a non-downstream process: induction, mental simulation, or some combination of the two. At t_{BEFORE}, we are entitled to take these processes for granted, ignoring possibilities in which Andy forms his belief by these processes and ENJOY is false. At t_{AFTER}, it is no longer presupposed that the process by which one acquires knowledge that ENJOY is a non-downstream process. In the case of knowledge about whether someone enjoyed a meal the previous day, it is typically presupposed that the relevant belief-forming process is a downstream process. Since a downstream process is presupposed as the relevant belief-forming process, different worlds comprise the set of relevant alternatives in assessing Andy's knowledge of ENJOY at t_{AFTER} as compared to t_{BEFORE}. Since induction and mental simulation are no longer presupposed to be the relevant processes, they are no longer taken for granted, so worlds in which Andy possesses the same inductive and mental simulation evidence and Beth did not enjoy the dish are no longer properly ignored. Andy's downstream evidence fails to distinguish between worlds in which she enjoyed the dish and worlds in which she did not. And so this explains why we are inclined to deny that Andy knows in this context and also why an assertion of (2) 'Yes, she liked it' would be inappropriate.

Easy foreknowledge cases fall out of a more general kind of context sensitivity due to differences in what reliable belief-forming processes are presupposed across contexts. To see this, consider the Cecilia's Departure case. In the first context, it is presupposed that Bianca's belief that Cecilia has departed is formed by a non-downstream process (some combination of induction and mental simulation). In such a context, Bianca counts as knowing that Cecilia has departed. In the second context in which Dieter calls Bianca and asks whether Cecilia has departed, it is not presupposed that Bianca's belief is formed by a non-downstream process. Given that beliefs about the past are typically formed by downstream processes, worlds in which Bianca has the same evidence, yet Cecilia has not departed become relevant and Bianca plausibly does not know the proposition in question and this explains why it seems inappropriate for her to flat-out assert that Cecilia has departed. Such a case is not an easy foreknowledge case since the event that the proposition is about (namely, Cecilia's departure) is earlier than both contexts. So the explanation for easy foreknowledge cases is not essentially temporal, but rather falls out of a more

general context sensitivity arising from presuppositions about the sort of process involved in forming the relevant belief.

Although the explanation for easy foreknowledge cases is not essentially temporal, the causal structure of our world, as well as our common knowledge of this causal structure, gives rise to such cases. Our beliefs about future events are not formed by downstream processes because there is no backwards causation from the event that our belief is about to our belief. Furthermore, either implicitly or explicitly, we are all aware of this fact. Suppose for a moment that seeing into the future was commonplace: a common and reliable way of forming beliefs about the future is by closing your eyes and pressing firmly on your eyelids with your index fingers. Suppose that doing this provides one with reliable images of what will happen. If this was a common and reliable way of forming beliefs about the future, and we were all aware that beliefs about future events are typically and reliably formed in this way, we may well deny that an individual who hasn't formed his beliefs in this way, but has done so rather on the basis of induction, knows the corresponding belief about the future. If we typically and reliably formed beliefs about the future in this way, and we are considering a context like t_{BEFORE} in which Andy has formed his belief that ENJOY by some combination of induction and mental simulation, and he has not formed it by pressing firmly on his eyelids, we may well deny that he knows ENJOY at t_{BEFORE} and judge that it would be inappropriate for him to assert (1). Worlds in which his belief is formed by some combination of induction and simulation, yet Beth does not enjoy the dish may well be relevant in such a context: the processes by which Andy formed his belief may not be taken for granted. But since in the actual world, beliefs about the future are not formed in this way, and induction and mental simulation are considered reliable ways of forming such beliefs, such processes *are* taken for granted and Andy counts as knowing ENJOY in t_{BEFORE}.

3.1 Objections and Replies

Like Ninan's account, the account I have proposed maintains that Andy knows ENJOY at t_{BEFORE} and fails to know it at t_{AFTER}. Fabrizio Cariani (2021) has recently provided some objections for "all epistemic accounts" of the Beth Case. i.e. accounts that explain the difference in assertibility between (1) and (2) in terms of a difference in Andy's knowledge of ENJOY between t_{BEFORE} and t_{AFTER}. The first objection derives from considering the following knowledge

claim at t_{AFTER}:

(5) Andy knew yesterday that Beth would enjoy the dish. (Cariani 2021, 268).

Cariani asks the proponent of epistemic accounts whether (5) is true. If it is true, he claims, the following "should sound good":

(6) Andy knew yesterday that Beth would enjoy the dish, but he doesn't know now that Beth did enjoy the dish, and the evidence on which he bases his belief hasn't changed.

Cariani claims that there is "something odd about [6], but epistemic accounts appear to sanction every component of it" (269). However, at least when levelled against a contextualist account of knowledge, this strikes me as a version of a well-worn objection; an objection that in my estimation has been adequately responded to by contextualists. Consider another well-known example motivating a contextualist account of knowledge: Keith DeRose's Bank Case (DeRose 1992). Keith and his partner drive to the bank on Friday to deposit a check and see a long line. Keith says to his partner "I know the bank is open tomorrow morning. Let's come back tomorrow." Suppose that the bank is in fact open tomorrow morning. Keith bases his belief on the fact that he went to the bank on a Saturday morning two weeks ago and it was open. The contextualist argues that in such a context it is plausible to claim that Keith counts as knowing that the bank will be open tomorrow morning. Suppose later on Friday, Keith's partner says to him "Do you know that the bank will be open tomorrow? Banks frequently change their hours." The contextualist claims that in such a context, it would be appropriate for Keith to reply, "I guess I don't know. Let me confirm". In such a context, Keith does not count as knowing that the bank is open tomorrow. At the later time on Friday when Keith's partner raises the possibility that the bank has changed its hours, we can consider the following claim:

(7) Keith knew earlier that the bank would be open tomorrow.

And we can ask the contextualist whether (7) is true. If it is true, then the following should sound good.

(8) Keith knew earlier that the bank would be open tomorrow, but he doesn't know now that the bank will be open tomorrow, and the evidence on which he bases his belief hasn't changed.

However, according to the contextualist, the meaning of 'knows' and 'knew' varies across contexts. 'know' in Keith's mouth at the earlier time has a different meaning from 'knew' in (7) evaluated at the later time. At the later time, more possibilities are relevant for evaluating whether 'Keith knows the bank is open tomorrow' is true and similarly for evaluating whether 'Keith knew earlier that the bank would be open tomorrow' is true. So the natural response for the contextualist is to deny that (7) is true at the later time. Although there are different mechanisms at play in the two examples (raising to salience vs presupposition change), in both the Beth Case and the Bank Case, the meaning of 'knows' (and 'knew') changes from the earlier context to the later context, and different possibilities are relevant for evaluating the truth of knowledge attributions. So the natural response for the contextualist to Cariani's objection is to deny that (5) is true at t_{AFTER}. And this seems in accordance with the account proposed. Evaluations of knowledge claims regarding ENJOY at t_{AFTER} presuppose that the relevant belief was formed by downstream processes and because neither Andy's belief at t_{BEFORE} nor his belief at t_{AFTER} was formed by downstream processes, 'Andy knows Beth enjoyed the dish' and 'Andy knew yesterday that Beth would enjoy the dish' are both false at t_{AFTER}. The contextualist should claim that at t_{AFTER} (5) and hence (6) are both false and so the contextualist is not committed to accepting the problematic (6) at t_{AFTER}.

Cariani's second objection, which he attributes to Matt Mandelkern and is also discussed by Ninan (2014, 304-5) in a different context, involves the following knowledge ascription at t_{AFTER}:

(9) Andy knows that Beth must have enjoyed the dish (Cariani 2021, 269).

Cariani claims that (9) is plausibly true at t_{AFTER}. If we grant that (9) is true at t_{AFTER}, "two big pieces of philosophical orthodoxy" commit us to claiming that Andy knows that Beth enjoyed the dish at t_{AFTER}. The first bit of philosophical orthodoxy is that *must*-claims entail their prejacents: *must*-ϕ entails ϕ. *Beth must have enjoyed the dish* entails that *Beth enjoyed the dish*. The second bit of philosophical orthodoxy is that knowledge is closed under entailment. The following argument can then be given for the claim that Andy knows ENJOY at t_{AFTER}:

1. At t_{AFTER} Andy knows that Beth must have enjoyed the dish.
2. *Beth must have enjoyed the dish* entails *Beth enjoyed the dish*. (*must*-ϕ

entails ϕ)

3. So at t_{AFTER} Andy knows that Beth enjoyed the dish. (1, 2, Closure).

Cariani concludes "if that conclusion is accepted, the assertibility asymmetry cannot be grounded in the epistemic asymmetry because there is no epistemic asymmetry" (269).

I think the proper response to this objection involves wading into some tricky issues regarding the nature of epistemic *must* claims. There is a puzzle associated with *must*-ϕ statements that has been dubbed by von Fintel and Gillies (2010) as "Karttunen's Puzzle" due to the fact that it was first explicitly stated by Karttunen (1972). The puzzle is that, on the one hand, *must* in English frequently expresses epistemic necessity. On the standard semantics for epistemic *must*-ϕ claims, ϕ is true in all possibilities compatible with what is known. And so for any ϕ, *must*-ϕ entails ϕ. In this sense, *must*-ϕ seems at least as strong as ϕ.

On the other hand, there is a clear intuition that *must*-ϕ claims express something *weaker* than ϕ claims. Consider the following two responses to the question 'Where are the keys?'

(10) The keys are in the drawer.

(11) The keys must be in the drawer.

Many share the intuition that (11) makes a weaker epistemic claim than (10). Hence, Karttunen's puzzle: the standard semantics for *must*-ϕ takes it to be at least as strong as ϕ, but assertions like (11) seem to assert something weaker than ϕ. Some have taken examples like (10) and (11) as evidence for the claim that *must*-ϕ is in fact weaker than ϕ.[14] Given the difference of opinion on whether *must*-ϕ is in fact weaker than ϕ, premise 2 of the above argument may not be supported by "philosophical orthodoxy". Accepting a weak semantics for *must*-ϕ would give us an easy way out.

Nonetheless, von Fintel and Gillies (2010) provide several compelling arguments against the view that *must*-ϕ is in fact weaker than ϕ , and I do think that one who accepts these arguments has resources for responding to Cariani's objection. Von Fintel and Gillies explain the contrast between (10) and (11) by positing a semantic presupposition as part of the lexical meaning of *must*. The

[14]Some who have provided a semantics for *must*-ϕ that is weaker than ϕ include Veltman (1985, see pp.161-163) and Kratzer (1991).

semantic presupposition is that the evidence for ϕ is indirect: *must-*ϕ carries with it an epistemic marker that the belief was formed on the basis of indirect evidence. However, it is unclear how the relevant notion of indirect evidence ought to be understood.[15]There is a further feature noted by Stone (1994) and Mandelkern (2019) that *must-*ϕ claims presuppose that is inferred from salient or publicly available premises. Stone paraphrases *must-*ϕ as 'some particular collection of facts A, salient in the common ground, provide a decisive reason to adopt the belief that []'(1994: 2) . Following Stone (1994), Mandelkern (2019) provides an attractive further development of the indirectness proposal in terms the following principle:

Support Assertions of *must-*ϕ are felicitous only if there is an argument for that is salient to the conversational participants.

Using the insights from von Fintel, Gillies, Stone, and Mandlekern, one can respond to the above argument by appealing to the pragmatic features at play in assertions of *must-*ϕ. Such a response involves denying premise 1 but explaining the appeal of the premise by granting that an assertion by Andy of:

(12) Beth must have enjoyed the dish

on Saturday morning may well be felicitous due to presupposition accommodation. Let me explain. If what one were to say would change the context, such as uttering 'I know I am not a brain-in-a-vat' in a low standards context, whether a knowledge attribution is true or false at a given context does not always line up with what would or would not be assertable. At t_{AFTER} Andy does not know must-ENJOY: that Beth must have enjoyed the dish. But an assertion of (12) on Saturday morning would change the context and may well be appropriate in the resulting context.

Both Stone (1994) and Mandelkern (2019) note that *must-*ϕ assertions are associated with certain patterns of accommodation concerning how the belief that was formed. Stone considers the following exchange:

(13) Ann asked, 'Where is the sugar?' and Mary answered, 'It must be in the cabinet over the fridge' (Stone 1994: 4)

[15]For one, 'indirect evidence' ought not be understood to include highly circuitous, causally downstream, belief-forming processes. Asserting 'Abraham Lincoln must have been shot' seems like an inappropriate way of expressing my belief (more appropriate for a detective arriving on the scene at Ford's Theater), even though there is a sense in which the evidence for my beleif is highly indirect, involving a long, convulted causal chain of perception, memory and testimony.

Stone notes, 'to make sense of Mary's answer in [13], one must assume that Mary has just seen something or figured something out from which she concludes that the sugar is in the cabinet over the fridge. Perhaps Mary has seen a telltale trail of white particles, or perhaps she has realized that only one cabinet remains in the kitchen which Ann has not ruled out' (1994: 4) In such cases 'it may be necessary to accommodate a salient argument with *must*' (1994). Along similar lines, Mandelkern gives an example in which his phone rings and he says:

(14) This must be my brother; let me take this.

Prior to the assertion (let us suppose), there is no argument salient to both speaker and hearer for the conclusion that his brother is calling. However, when the assertion occurs, the hearer accommodates the relevant presupposition: that the speaker has formed his belief on the basis of an argument for the conclusion that his brother is calling: his brother said he would call around 2pm and it is around 2pm, etc. Similarly, if Andy were to assert, 'Beth must have enjoyed the dish', Andy's conversational partners would accommodate the presupposition that he has formed his belief that Beth enjoyed the dish on the basis of an argument for this conclusion: likely an inductive one similar to the one on which he based his belief at t_{BEFORE} he has made the dish several times in the past, Beth has always liked it, etc. Given that an assertion of (12) involves accommodating the presupposition that Andy believes ENJOY on the basis of an argument salient to the conversational participants, the context resulting from an assertion of (12) is different from the context considered in the Beth Case at t_{AFTER}, and in accommodating Andy's assertion, presuppositions about the relevant belief-forming process have changed. In conclusion, at t_{AFTER} Andy fails to know that Beth must have enjoyed the dish, just as he fails to know that she enjoyed the dish; however, were he to assert on Saturday morning, 'Beth must have enjoyed the dish', conversational participants may well accommodate the presupposition that Andy believes ENJOY on the basis of a salient argument and, given the new context, his assertion may well be felicitous.

4 Conclusion

I have highlighted fundamentally different kinds of reliable processes by which we acquire beliefs about past, present and future events. Furthermore, I have argued that we are sensitive to the kind of process by which such beliefs are formed in our attributions of knowledge and judgments about the appropriateness of

assertions. I have claimed that these pragmatic presuppositions about which belief-forming process is in play leads to a new kind of context-sensitivity with respect to knowledge attributions. The framework of context-sensitivity naturally extends to Ninan's easy foreknowledge puzzle, explaining our judgments about the appropriateness of (1) and (2) without making the phenomenon essentially temporal. It is perhaps no surprise that in assessing whether a subject knows a proposition concerning the future, or whether assertions about the future are appropriate, we typically apply different considerations from those involved in assessing knowledge and assertions about the past. Afterall, the *way* by which we gain knowledge of the future is of a fundamentally different sort than the way by which we typically gain knowledge of past events. And it is no surprise that this deep difference works its way into our assessments of knowledge and assertion.[16]

References

[1] G. E. M. Anscombe. *Intention*. Harvard University Press, 2000.

[2] Sara Aronowitz and Tania Lombrozo. Learning Through Simulation. *Philosophers' Imprint*, 20, 2020.

[3] Peter W. Battaglia, Jessica B. Hamrick, and Joshua B. Tenenbaum. Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45):18327–18332, November 2013. Publisher: National Academy of Sciences Section: Biological Sciences.

[4] Michael Blome-Tillmann. *Knowledge and Presuppositions*. Oxford University Press, Oxford, 2014.

[5] Fabrizio Cariani. *The Modal Future: A theory of future-directed thought and talk*. Cambridge University Press, 2021.

[6] Keith DeRose. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4):913–929, 1992. Publisher: Wiley-Blackwell.

[16] I would like to thank Mahrad Almotahari, Phillip Bricker, Jesper Kallestrup, Hilary Kornblith, Federico Luzzi, Bailie Peterson, Patrick Todd, Brandt van der Gaast, Jacek Wawer, the audience of the Utrecht University Philosophy Colloquium, and the audience of BrickerFest 2022 for helpful comments on earlier versions of this paper. Additionally, the comments from the editor and two anonymous referees for this journal were very helpful and substantially improved this paper.

[7] Alvin I. Goldman. A Causal Theory of Knowing. *The Journal of Philosophy*, 64(12):357–372, 1967. Publisher: Journal of Philosophy, Inc.

[8] Alvin I. Goldman. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Philosophy of Mind Series. Oxford University Press, New York, 2006.

[9] Jessica B Hamrick. Analogues of mental simulation and imagination in deep learning. *Current Opinion in Behavioral Sciences*, 29:8–16, October 2019.

[10] Lauri Karttunen. Possible and Must. *Syntax and Semantics Volume 1*, pages 1–20, December 1972. Publisher: Brill.

[11] Angelika Kratzer. Modality. pages 639–650, 1991.

[12] David Lewis. Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.

[13] Matthew Mandelkern. What 'must' adds. *Linguistics and Philosophy*, 42(3):225–266, June 2019.

[14] Richard Moran. Anscombe on 'Practical Knowledge'. *Royal Institute of Philosophy Supplements*, 55:43–68, September 2004. Publisher: Cambridge University Press.

[15] Shaun Nichols and Stephen P. Stich. *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding Other Minds*. Oxford Cognitive Science Series. Oxford University Press, Oxford, 2003.

[16] Dilip Ninan. Taste predicates and the acquaintance inference. *Semantics and Linguistic Theory*, 24(0):290–309, August 2014. Number: 0.

[17] Dilip Ninan. Assertion, Evidence, and the Future. *Philosophical Review*, page 52, 2022.

[18] Daniel L. Schwartz and Tamara Black. Inferences through imagined actions: Knowing by simulated doing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(1):116–136, 1999. Place: US Publisher: American Psychological Association.

[19] Kieran Setiya. Anscombe on Practical Knowledge. In *Practical Knowledge: Selected Essays*. Oxford University Press, 2016.

[20] Martin Smith. What Else Justification Could Be. *NoÃ»s*, 44(1):10–31, 2010. Publisher: Wiley.

[21] Robert Stalnaker. Pragmatic Presuppositions. In Robert Stalnaker, editor, *Context and Content*, pages 47–62. Oxford University Press, 1974.

[22] Robert Stalnaker. Assertion. *Syntax and Semantics (New York Academic Press)*, 9:315–332, 1978. Publisher: New York: New York Academic Press.

[23] Robert C. Stalnaker. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford Cognitive Science Series. Oxford University Press, Oxford, 1999.

[24] Robert C. Stalnaker. Common Ground. *Linguistics and Philosophy*, 25(5-6):701–721, 2002. Publisher: Springer Netherlands.

[25] Matthew Stone. The reference argument of epistemic must. *Proceedings of IWCS*, 1:1–10, 1994.

[26] Frank Veltman. *Logics for Conditionals*. PhD thesis, January 1985.

[27] Kai von Fintel and Anthony S. Gillies. Must . . . stay . . . strong! *Natural Language Semantics*, 18(4):351–383, December 2010.