

# Easy Foreknowledge

Dilip Ninan, Tufts University  
dilip.ninan@tufts.edu  
Zermatt Philosophy Workshop  
February 2019

## 1. Introduction

In the Western philosophical tradition, there are two well-known obstacles to knowing the future:

- Aristotle's problem of future contingents (*De Interpretatione ix*):

Unlike the past and present, the future is open. Suppose that there are possible futures in which a sea battle takes place tomorrow and possible futures in which no sea battle takes place tomorrow. If there is no sense in which one of the futures is the privileged, actual future, then it is presumably neither true nor false that there will be a sea battle tomorrow, and neither true nor false that there will not be a sea battle tomorrow. If one can only know what is true, then one cannot know whether or not there will be a sea battle tomorrow.

- Hume's problem of induction:

Our present evidence is consistent with the future's taking some radically unexpected course.

The contrary of every matter of fact is still possible, because it can never imply a contradiction... That the sun will not rise tomorrow is no less intelligible a proposition, and implies no more contradiction, than the affirmation, that it will rise. We should in vain, therefore, attempt to demonstrate its falsehood. (*An Enquiry Concerning Human Understanding*, §4.2)

Even if there is only one *actual* future in which the sun rises tomorrow, one might worry that our present evidence doesn't rule out the epistemic possibility of it's not rising. Perhaps there are multiple possible worlds which are qualitatively identical to the actual world up until the present moment. Since some of these are worlds in which the sun does not rise tomorrow (we can suppose), how can one *know* that the sun will rise tomorrow? After all, one's evidence would be exactly the same were one in one of these worlds with a cold, dark future.

Comment:

- Upshot of these arguments: knowledge of the future is more difficult to obtain than knowledge of the past and present.
- It is somewhat surprising then to learn that there are cases that pull in the opposite direction, cases in which it appears *easier* to obtain knowledge of the future than to obtain knowledge of the past or present.
- Arguably, the following are two such cases.

### Beth case

Andy is a personal chef to a wealthy entrepreneur, Beth. Andy is making a new dish for Beth's dinner tonight. Based on his knowledge of the sorts of foods that Beth usually likes, Andy says to his friend Chris, *Beth will love 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 day, Andy learns that Beth ate the dinner he had made for her, but he doesn't hear specifically whether or not she liked it. One of Andy's friends asks Andy, *Did Beth enjoy the dish you made for her yesterday?*

I think it would seem odd here for Andy to flat-out assert that she did, i.e. to say, *Oh, yes, she loved 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 enjoyed the dish. Absent evidence of that sort, it would be better for Andy to hedge in some way, i.e. to say, *She probably loved it* or *She must have loved it – it had all her favorite ingredients.*

Beth case:

Friday 4pm (*F4*): Andy can say, "Beth is going to like this when she eats it at 7pm."

(For the sake of ease, we stipulate that Andy knows that Beth is going to eat the dish at 7pm.)

Friday 7pm (*F7*): Beth eats and enjoys the dish in question.

Saturday 10am (*S10*): Andy can't say, "Beth liked the dish when she ate it at 7pm."

### Rain case

Ellen has been visiting her friend Frank in Chicago for the last few days, but the visit is over and he is driving her to the airport. He's telling Ellen about a big outdoor concert he's planning to attend this weekend, but he's worried about the weather. He asks Ellen if she can check the forecast. Ellen looks on her phone, and says, *Bad news – it's going to rain all weekend.* Frank replies, *Oh, that's too bad.* Ellen catches her flight back to Boston. It does indeed rain all weekend in Chicago, and Frank's concert gets cancelled.

On Monday, Ellen goes to work, and bumps into a co-worker, who is also a friend of Frank's. The co-worker also knew about Frank's plan to attend the concert, but he hasn't yet heard whether or not it was cancelled. He asks Ellen what the weather in Chicago was like over the weekend. Ellen hasn't heard from Frank or anyone else what the weekend in Chicago was like. Again, it seems to me that it would be inappropriate for Ellen to flat-out assert that it rained all weekend in Chicago, and this is so even if she makes it clear what her evidence for that assertion is. If she wants to make a comment about what the weather was like in Chicago over the weekend, she needs to make a hedged assertion, e.g. *It was supposed to rain all weekend.*

Rain case:

Friday: Ellen can say, "It's going to rain in Chicago on Saturday."

Saturday: It rains in Chicago.

Monday: Ellen can't say, "It rained in Chicago on Saturday."

Two comments:

- (i) I'm moved from talking about knowledge to talking about assertion. But the connections between these notions will soon emerge.
- (ii) What this shows:

*There are cases in which: (i) one can assert it will be the case that  $\phi$  at time  $t$ , (ii) one's relevant evidence does not change between  $t$  and later  $t'$ , and yet (iii) one cannot assert it was the case that  $\phi$  at  $t'$ .*

This does not show:

*For any* case in which (i) and (ii) obtain, (iii) obtains.

This universally quantified claim is false, as we will discuss.

**Question:** Why is Andy in a position to say, “Beth is going to like this” while he’s preparing the dish, but not in a position to say, “Beth liked that dish” the next day, despite the fact that he didn’t lose or gain any relevant evidence between the two times?

I’m going to consider two possible answers to our question:

Hypothesis (A): The explanation of this phenomenon is rooted in a **semantic** difference between future operators, on the one hand, and the present and past tenses, on the other.

Future operators have a certain epistemic content to their meaning that the present and past tenses lack, and it is this difference that explains our observations.

Hypothesis (B): The explanation of this phenomenon is essentially **pragmatic/epistemological**.

As we move through time, more (or better) evidence becomes available. The evidence *needed* to make assertion at a given time  $t$  is a function of the evidence *available* at  $t$ . Since more/better evidence becomes available as we move through time, it is generally ‘harder’ to assert *it was the case that  $\phi$*  at later time  $t'$  than it is to assert *it will be the case that  $\phi$*  at earlier time  $t$ .

Plan:

- I’m going to start by considering one version of Hypothesis (A), and then I’ll turn to consider a version of Hypothesis (B).
- Of the hypotheses I discuss, Hypothesis (B) seems more likely to be true. As I develop it, Hypothesis (B) commits me to some claims about knowledge, which I will spend some time exploring at the end of my talk.

Three assumptions:

(i) Eternal propositions:

- Since I want to compare what an agent knows or asserts at one time with what she knows or asserts at a later time, it is easier to conduct the discussion assuming that propositions do not change their truth value over time: they are *tenseless* or *eternal*.
- I don’t think anything really hangs on this. *Temporalists* (philosophers who think that propositions vary in truth value over time) will have to give us some sort of story about what (e.g.) it is to retain knowledge or to change one’s mind, and once they tell that story, we could likely conduct our discussion in temporalist terms.

(ii) Knowledge norm:

- For simplicity, I’m going to assume that knowledge is the norm of assertion (Williamson, 1996). According to this view, knowledge is governed by the following norm: one must assert  $\phi$  only if one knows  $\phi$ . I take this to be one way of interpreting Grice’s Maxim of Quality.
- This is mostly for convenience (though I myself do accept this doctrine). Most of what I say could be re-cast in terms of the notion of *warranted assertability* or *being in a position to assert*.

(iii) Past tense:

- On Saturday at 10am, the sentence “Beth liked the dish at 7pm” carries two pieces of information:
  - the (tenseless) proposition that Beth likes the dish at 7pm on Friday, and

- the proposition that Friday 7pm is earlier than the utterance time (Saturday 10am).
- I’m going to mostly ignore the second piece of information, and focus on the first. Call this proposition “ $\beta$ ,” and represent it as follows:

$$\beta := \{w : \text{Beth likes the dish at 7pm on Friday in } w\}$$

## 2. Future operators as epistemic modals

Hypothesis A says that future operators have a certain epistemic content.

Some evidence that *will* is an epistemic modal:

- Epistemic (i.e. non-future) uses of *will*. (Palmer 1986, Enç 1997, Winans 2016)
  - (1) It’s 5pm, so John will be in Paris now.
- Like other modals, *will* participates in modal subordination. (Roberts 1989, Klecha 2013, Cariani and Santorio 2017)
  - (a) If the supplies arrive tomorrow, it will be late in the day. They will contain three boxes of cereal.
  - (b) # If the supplies arrived yesterday, it was late in the day. They contained three boxes of cereal.
- Like epistemic modals, *will* obviates the ‘acquaintance inference’ on taste predicates. (Klecha 2013, Ninan 2014)
  - (2) Those cookies are delicious.   *suggests the speaker has tasted them*
  - (3) Those cookies are probably/must be delicious.   *no such suggestion*
  - (4) Those cookies will be delicious.   *no such suggestion*

Proposal:

- There’s a type of proposal for what the epistemic content of *will* is that, if true, would help with our puzzle.
- According to this type of proposal, *will* quantifies over all the *most likely future possibilities* (Kaufmann 2005) or over all the *most normal future possibilities* (Copley 2009).
- But – and here is the crucial feature – these proposals don’t require that the actual future is among the most likely or most normal future possibilities.
- So, when Andy says at 4pm on Friday, “Beth will like this,” he is simply saying that in all the most normal/likely future possibilities, Beth likes the dish.
- This can be true even if, in the actual world, Beth doesn’t like the dish, since the actual world might not be sufficiently likely or normal.
- On this proposal, the past tense – unlike the future tense – has no epistemic content to it. So the sentence “Beth liked the dish” when considered on Saturday morning, is true in the actual world only if Beth liked the dish in the actual world.
- So the proposition Andy asserts when he says, on Friday afternoon, “Beth will like the dish,” does not entail the proposition that would be expressed by the past tense sentence “Beth liked the dish” were that sentence uttered on Saturday morning.
- That opens up the possibility that Andy knows the first proposition, but not the second, which would help to explain why he’s in a position to assert the first, but not the second.

Future possibilities:

- Given a world  $w$  and time  $t$ ,  $w'$  is a *future possibility with respect to  $w$  and  $t$*  just in case  $w'$  is just like  $w$  up until  $t$ .  
See Lewis (1979) on the notion of one world's being just like another up until a time  $t$  (they 'match perfectly in matters of particular fact').
- If  $F_{w,t}$  is the set of future possibilities w.r.t.  $w$  at  $t$ , then let  $N(F_{w,t})$  be the subset of  $F_{w,t}$  that includes all and only the the future possibilities that are most *likely* (Kaufmann 2005) or most *normal* (Copley 2009).
- Note that world  $w$  need not be an element of  $N(F_{w,t})$  since since it might be that something abnormal or unlikely happens at time  $t$  in world  $w$ .

*Beth will like the dish at 7pm:*

- On this approach, at 4pm on Friday, the sentence "Beth will like the dish at 7pm," carries two pieces of information:
  - the proposition that in *all the most normal/likely post-F4 future possibilities*, Beth likes the dish at 7pm on Friday.
  - the proposition that Friday 7pm is no earlier than the utterance time (Friday 4pm).
- Let's call this first proposition " $\alpha$ " and represent it as follows:

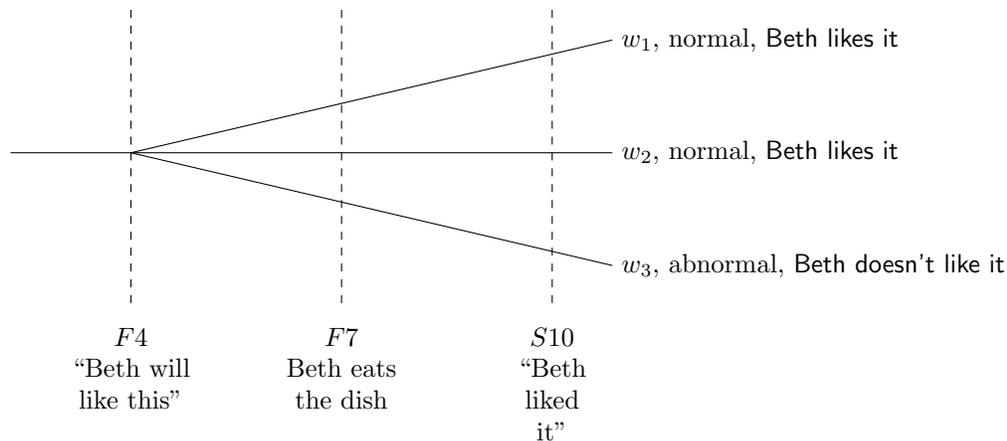
$$\alpha := \{w : \forall w' \in N(F_{w,F4}): \text{Beth likes the dish at 7pm on Friday in } w'\}$$

- In our puzzle case, given Andy's knowledge of Beth's tastes, we can assume that he knows at  $F4$  that, in all the most normal/likely future possibilities, Beth likes the dish.  
That is, on Friday afternoon, Andy knows  $\alpha$ , and so is in a position to assert it.

*Beth liked the dish at 7pm:*

- Now, according to this proposal, on Saturday morning, "Beth liked the dish at 7pm" expresses (*inter alia*) proposition  $\beta$ :
$$\beta = \{w : \text{Beth likes the dish at } F7 \text{ in } w\}$$
- It's consistent with what we've said so far that Andy doesn't know this proposition on Saturday morning, even if he retain his knowledge of  $\alpha$  as he moves from Friday afternoon to Saturday morning.
- This is because, for all Andy knows, his world might be one in which the post- $F4$  future is an abnormal one in which Beth fails to like the dish.
- So we might suppose that Andy doesn't know  $\beta$  on Saturday morning, and that this explains why he's not in a position to assert it.

Diagram:



Suppose Andy's world is  $w_1$ , and  $\{w_1, w_2, w_3\}$  are the worlds compatible with what Andy knows.

Key features of this account:

- *Beth will like the dish at 7pm* expresses  $\alpha$  at on Friday afternoon.
- *Beth liked the dish at 7pm* expresses  $\beta$  on Saturday morning.
- Proposition  $\alpha$  does not entail proposition  $\beta$ .
- At all times, Andy knows  $\alpha$  but doesn't know  $\beta$ . So Andy never gains nor loses knowledge as he moves through time.

Problems with this account:

- This approach solves our puzzle, but faces at least two problems.
- First problem:  
The account appears to predict that *will* fails to commute with negation, but the evidence suggests otherwise (MacFarlane 2014, Cariani and Santorio 2017).  $\forall \neg \neq \neg \forall$ 
  - But one could stipulate that *will* always takes wide-scope over negation (Winans 2016) or try other maneuvers to avoid this prediction.
- Second problem:  
The account seems to predict the wrong truth-conditions for *it will be the case that  $\phi$*  (Cariani and Santorio 2017). I expand on this below.

Truth-conditions:

- Imagine that we are about to flip a fair coin. We make the following bet: if (5) is true now, I'll pay you \$10; if (5) is false now, you pay me \$10.
    - (5) The coin will land heads.
- We flip the coin. It lands heads. So it looks like I have to pay up.
- But the coin was fair. So it wasn't true at the time of the bet that in all the most normal/likely future possibilities, the coin would land heads. In some of the normal/likely future possibilities, the coin lands tails.
- So (5) was false when you uttered it, and you owe me \$10.

- Actually, why even flip? If we know the coin is fair, then we know, prior to flipping it, that it is not the case that in all the most normal/likely future possibilities the coin lands heads.  
So we can skip the coin toss, and you can just pay me \$10.
- This also yields incorrect predictions about how confident we should be, prior to the flip, that (5) is true.

Summary:

- Hypothesis (A) was the broad claim that the source of our puzzle lies in the fact that future operators have a certain epistemic content to them.
- There is still evidence for that (see above).
- All I've done here is examined one way of trying to understand what that epistemic content is. I've argued that, although the resulting approach would help with our puzzle, that approach is implausible for independent reasons.
- But perhaps the epistemic content of future claims is not best understood as an aspect of their truth conditions; perhaps it is part of the 'not-at-issue' content of future claims. I won't explore this idea further here, but perhaps it is worth exploring.

### 3. A pragmatic-epistemic account

Simple semantics for future operators:

- On Friday at 4pm, the sentence “Beth will like the dish when she eats it at 7pm” carries two pieces of information:
  - the proposition that Beth likes the dish at 7pm on Friday, and
  - the proposition that Friday 7pm is later than the utterance time (Friday 4pm).
- The first proposition is just  $\beta$ .
- $\beta = \{w : \text{Beth likes the dish at 7pm on Friday in } w\}$   
So on this approach, “Beth will like the dish” at  $F4$  and “Beth liked the dish” at  $S10$  are equivalent, *modulo* the information they contain about how the utterance time relates to the time of Beth's eating.
- Abstracting away from this temporal information, both utterances express proposition  $\beta$ , the proposition that Beth likes the dish at 7pm on Friday.
- Then the question becomes: Why is Andy in a position to assert  $\beta$  on Friday afternoon but not on Saturday morning?

Hypothesis (B):

- The reason Andy is in a position assert  $\beta$  on Friday afternoon is that he knows it on Friday afternoon. The reason Andy is not in a position to assert  $\beta$  on Saturday morning is that he does not know it on Saturday morning.
- On this approach, Andy loses knowledge despite not losing or gaining any relevant evidence. He loses knowledge simply by moving through time.  
Why is this? Why he lose knowledge of  $\beta$  as he moves from Friday afternoon to Saturday morning?
- Although Andy's relevant evidence doesn't change between Friday afternoon and Saturday morning, the evidence that is *in principle available* to Andy changes between those two times.

- Note that, on Saturday, Andy is in a position to get *more evidence* that bears on the question of whether Beth liked the dish. For example, he could call Beth up and ask her if she liked the dish. And she might then tell him that she did indeed like it.
- Note also that, on Saturday, Andy is in a position to get evidence *of a different kind*.  
Beth's telling Andy that she liked the dish is evidence that is *directly connected* to the fact that Beth liked the dish. It is connected to the underlying fact in a way that the evidence available to him on Friday afternoon is not.
- I'll come back to this issue of quantity of evidence vs. kind of evidence in a moment, but first I want to pause to make a few comments.

Three comments:

(i) Knowledge and evidence

- Here is what looks like a truism:

Knowledge-Evidence Link

$S$  knows  $p$  at time  $t$  only if  $S$ 's total evidence at  $t$  supports  $p$ .

- Given this truism, it is natural to think that one can lose knowledge by losing evidence. But can also lose knowledge because one *gains* misleading evidence. Here is a case from Williamson (2000).

Suppose at time  $t_1$ , I see a red ball and a black ball put into an empty bag. Thus, I come to know at that time that there are two balls in the bag, one red and the other black. I then see that in the first thousand draws with replacement, a red ball is drawn every time. This evidence, though compatible with what I earlier knew, should give me strong reason to doubt that there is a black ball in the bag. After these draws have been made, say at  $t_2$ , my total evidence will no longer support the claim that, of the two balls in the bag, one is black.

- What's interesting about our cases is that one loses knowledge despite no obvious change in one's relevant evidence.

(ii) Knowledge doesn't supervene on one's evidence plus practical stakes. The supervenience base for knowledge must include facts about what evidence is available to the subject.

(iii) Connection to Harman's cases (Harman, 1973).

*Assassination.* A political leader is assassinated. An enterprising and reliable journalist witnesses the assassination and writes a report on it, which is then printed in the final edition of a newspaper. Jill reads the newspaper report, and comes to believe that the political leader is dead. But later the associates of the political leader, fearing a coup, decide to pretend that the bullet hit someone else. On nationwide television they announce that an assassination attempt has failed to kill the leader but has killed a secret service man by mistake. Rationally trusting this announcement, most people come to abandon their earlier beliefs about the death of the political leader. But Jill wasn't watching television when this announcement was broadcast. She continues to believe that the leader was assassinated.

What do Harman's cases show?

- Harman thinks that, at the end of the first story, Jill does not know that the leader was assassinated.
- Let's suppose that, *prior* to the television announcement that Jill did not hear, Jill knew that the leader was assassinated.
- Then if Harman's judgment that Jill does not know this at the end of the case, Jill also loses knowledge despite not losing or gaining any relevant evidence.

- As Harman describes these cases, knowledge is undermined by ‘evidence one does not possess.’
- So Harman’s case also shows that knowledge doesn’t supervene on one’s evidence plus the stakes. And his case also seems to confirm our suggestion that what counts as sufficient evidence for  $S$  to know  $p$  at  $t$  depends on what evidence is *available* to  $S$  at  $t$ .
- One difference between our cases and Harman’s is that, in our cases, the subject becomes aware that more evidence that bears on the matter has become available; in Harman’s cases, the subject does not become aware in this way.

Two principles:

- Knowledge-Available Evidence Link 1

$S$  knows  $p$  at time  $t$  only if  $S$ ’s evidence at  $t$  supports  $p$ , and  $S$  has *enough* of the evidence relevant to  $p$  that is available to  $S$  at  $t$ .

- Knowledge-Available Evidence Link 2

$S$  knows  $p$  at time  $t$  only if  $S$ ’s evidence at  $t$  supports  $p$ , and  $S$  has the right *kind* of evidence relevant to  $p$  that is available to  $S$  at  $t$ .

**Question:** On Saturday, does Andy fail to know that Beth liked the dish because (i) he doesn’t have enough evidence, (ii) he doesn’t have the right kind of evidence, or (iii) both?

What does Andy lack on Saturday?

- It seems that in order for Andy to know on Saturday that Beth liked the dish, he needs more evidence.
- But is it just that he doesn’t have *enough* evidence? Could Andy come to know that Beth liked the dish by finding out even more about Beth’s taste in food? Perhaps.
- But it’s tempting to think that Andy needs evidence of a different *kind* from the evidence he currently possesses.
- What kind does he need? Well, we noted that it would suffice if Beth simply told Andy that she liked the dish in question.
- One thing about that evidence is that it is *causally downstream* from the fact that Beth liked the dish. It is an *effect* of Beth’s liking the dish .
- Andy’s evidence on Friday afternoon is *causally upstream* from the fact that Beth liked the dish. His evidence consists of *causes* of Beth’s liking the dish.

Knowledge and causation?

- Knowledge-Causation Link

$S$  knows  $p$  at time  $t$  only if:

if it is possible for  $S$  to have evidence for  $p$  at  $t$  that is caused by  $p$ , then  $S$  has such evidence for  $p$  at  $t$ .

- On Friday afternoon, it isn’t possible for Andy to have evidence for the claim that Beth likes the dish at 7pm on Friday that is caused by her liking it (no backwards causation). So this principle doesn’t require Andy to have such evidence.
- But on Saturday morning, it *is* possible for Andy to have such evidence. So the link principle says that he must have such evidence in order to know that Beth likes the dish at 7pm on Friday. Since he lacks such evidence on Saturday, the principle entails that he doesn’t know, on Saturday morning, that Beth likes the dish at 7pm on Friday.

The problem:

- The only problem with this is that the Knowledge-Causation Link principle is almost certainly false.
- **Jack case:**

Jack tells me on Friday that he's going to Portland on Saturday for two weeks. On Sunday, I bump into Jill, a mutual friend. She asks me how Jack's doing. I say, "He's great. He's in Portland right now for a couple of weeks."

- In this case, I know that Jack is in Portland (or so it seems to me). But I don't have any evidence that is caused by Jack's being in Portland, even though it is possible for me to get such evidence (I could just call Jack up).

This point is important for another reason:

- Shows that there are cases in which one can assert *it will be the case that  $\phi$  at  $t$* , one neither gains nor loses any relevant evidence between  $t$  and later time  $t'$ , and one can assert *it was the case that  $\phi$  at  $t'$* .
- What separates the **Jack case** from the **Beth case**? Is there a principle that will explain why these two cases differ?
- I have to admit that I'm a bit skeptical about finding a principle that will divide these cases correctly.

**Concluding remarks:**

- Don't let my failure to answer this question distract from the broader pragmatic-epistemic hypothesis, which (a) does seem to help explain what's going on with the Beth case, and (b) avoids the problems facing the version of Hypothesis (A) that we considered.
- That said, I want to reiterate a point I made earlier, which is that my arguments against Hypothesis (A) – the semantic hypothesis – were really only arguments against a particular version of that hypothesis.
- Is it just an accident that future operators can be used as epistemic modals? How should we account for the facts about *will* that suggested it was some sort of epistemic modal?

## References

- Cariani, F. and Santorio, P. (2017). *Will done better: Selection semantics, future credence, and indeterminacy*. *Mind*, **127**(505), 129–165.
- Copley, B. (2009). *The semantics of the future*. Routledge.
- Enç, M. (1997). Tense and modality. In S. Lappin, editor, *The Handbook of Contemporary Semantic Theory*, chapter 13. Blackwell, Oxford.
- Harman, G. (1973). *Thought*. Princeton University Press, Princeton.
- Kaufmann, S. (2005). Conditional truth and future reference. *Journal of Semantics*, **22**(3), 231–280.
- Klecha, P. (2013). Diagnosing modality in predictive expressions. *Journal of Semantics*, **31**(3), 443–455.
- Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*, **13**, 455–476.
- MacFarlane, J. (2014). *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford.
- Ninan, D. (2014). Taste predicates and the acquaintance inference. In *Semantics and Linguistic Theory*, volume 24, pages 290–309.
- Palmer, F. (1986). *Mood and Modality*. Cambridge University Press, Cambridge, UK.
- Roberts, C. (1989). Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy*, **12**(6), 683–721.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, **105**(4), 489–523.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press, Oxford.
- Winans, L. (2016). *Inferences of Will*. Ph.D. thesis, UCLA.