Chapter 5
Relevance in Epistemic Modal Disagreement

Jesse Fitts

Abstract I argue that pragmatic considerations explain puzzling epistemic modal disagreement cases. In particular, I claim that there are two different types of information sources involved in epistemically modalized propositions. One information source is a first-person epistemic state, or a group of epistemic states; another is a third-person, external source of information. This distinction helps make sense of felicitous and infelicitous responses in epistemic modal disagreement cases, which I go through in some detail.

Keywords Epistemic modality · Disagreement · Retraction · Contextualism · Relativism

5.1 Introduction

According to a simplified version of the orthodox theory of epistemic modals, *contextualism*,[1] modal particles quantify over domains that context usually supplies or otherwise an adverbial modifier supplies. In a *bare epistemic modal* (BEM), e.g., "Jones might be in his office," context supplies some knowledge that determines a domain over which the quantifier supplied by "might" ranges over and is true just in case that knowledge leaves open the truth of the *prejacent*, i.e., the unmodalized

[1] See, among many other papers, Kratzer (2012b/1981), DeRose (1991), Hawthorne (2007), Cappelen and Hawthorne (2009), Dowell (2011), Schaffer (2011), Dorr and Hawthorne (2013). For a contextualist view that departs in some ways from orthodoxy, see the "cloudy contextualism" of Fintel and Gillies (2011).

J. Fitts (✉)
Department of Philosophy, University of Nevada, Las Vegas, NV, USA
e-mail: jesse.fitts@unlv.edu

D. Zeman, M. Hîncu (eds.), *Retraction Matters*, Synthese Library 488,
https://doi.org/10.1007/978-3-031-66081-8_5

proposition that Jones is in his office.[2] Various authors have argued against orthodoxy based on modal disagreement cases in which a speaker's epistemically modalized assertion is challenged, and she should *retract* her initial assertion.[3] As the literature developed, the takeaway from these cases became increasingly complicated. One complication is that retracting a challenged epistemic modal is not always felicitous—we have a spread of cases in which retraction is felicitous and others where *digging in*, or refusing to retract, is felicitous. Another complication is that when someone disagrees with an epistemically modalized assertion, it's unclear exactly what she disagrees with—the entire epistemic modal or only the prejacent. On this point, Fintel and Gillies (2008, p. 83) note, "[b]efore resorting to the CIA's [the relativist's] extraordinary measures, we would like to see it carefully argued that flexibility in the target of denials and acceptances does not explain what needs explaining"—this paper partially takes up this challenge.

I have two goals in this paper. The first is to highlight and analyze a set of distinctions that are perhaps underappreciated in the epistemic modals literature. As the name suggests, epistemic modals involve epistemic content of some kind, typically knowledge. However, epistemic modals can sometimes involve, as many linguists and sometimes philosophers recognize, external information from, e.g., books, newspapers, etc. The first distinction highlights this difference in information type, which I label *first-* and *third-person*. The second distinction concerns conversational goals. Sometimes conversations revolve around locating the actual world in logical space; other times, around filling out logical space with no special regard for the actual world's location. I first explore these distinctions, which, while interesting in their own right, relate to the second goal of this paper.

My second goal will be to argue that conversational goals and information type can explain some retraction and digging-in intuitions. Roughly, I argue that first-person information sources are epistemic information sources that tend to pair with the conversational goal of locating the actual world in logical space. Third-person information sources are static information sources that tend to pair with conversational goals that sometimes don't involve locating the actual world in logical space. In retraction cases, I argue that certain assertions' irrelevance to the discourse can explain some of our felicity intuitions. I explain irrelevance in terms of a mismatch in the participants' conversational goals, which pair with information type.[4]

[2] See Kratzer (2012a/1977, 2012b/1981) for the basic logical structure of modals involving quantifiers, a (restricted) domain, and the prejacent.

[3] The retraction data was an initial motivation for relativism, the view that the content of an epistemically modalized sentence's truth value varies across contexts of assessment. For relativist papers in which retraction played an especially important role, see MacFarlane (2003, 2009) and Stephenson (2007a, b). For other relativist papers, see Egan et al. (2005) and Egan (2007). See Marques (2018) for an argument that retraction intuitions don't support relativism. More recently, retraction (among other linguistic intuitions) has been empirically studied to both contextualist and relativist ends, which I discuss below. See, e.g., Knobe and Yalcin (2014), Khoo (2015), Beddor and Egan (2018), and Kneer (2022).

[4] Outline of paper: §5.2 distinguishes first- and third-person information sources and explains conversational purpose, §5.3 focuses on the most relevant and difficult disagreement cases, and §5.4

5.2 First–Third-Person Distinction and Conversational Purpose

This section introduces and expounds on the two explanatory distinctions, starting with the first–third-person distinction followed by conversational purpose.

5.2.1 The First–Third-Person Distinction

Different kinds of information can be involved in epistemic sentences. The most common distinction, starting with the early work of Angelika Kratzer,[5] is between whether the information involved in an epistemic sentence is factual or not. The terminology in this area can be confusing.[6] Let's settle on distinguishing *realistic* information from *non-realistic* information. The idea is that the set of worlds that is quantified over will be determined by information that is true at the evaluation world—realistic—or not—non-realistic.

This distinction is closely related to another (as far as I know) novel distinction: between first-person and third-person takes on information. Suppose I say the following: "Jones might be home. His light is on." In such a situation, I'm not sure, or we're not sure, whether Jones is home. What I or we know doesn't rule it out, so I say that Jones might be home. The truth conditions, according to a simplified contextualist analysis, are that the information source, in this case *knowledge*, leaves open the prejacent—the proposition that Jones is in his office (at the time in question). Let me highlight some features of the information source in this case to contrast with another shortly. First, the information source is epistemic: knowledge. Second, the information source is *dynamic*. What I mean by this is that the information source is part of a constantly changing system, which I update, in the *epistemic sense*, as I find my way through the world. Third, I have a kind of commitment to the information source such that it doesn't make sense to say that I could disbelieve or be indifferent to it. Fourth, the information source is *perspectival* because I consider it part of my perspective.

There is a very different kind of epistemic-modal-involving situation. Let's say a while ago, I opened up the *Fox News* website. Various articles claim that the available information leaves open whether President Obama was born in the United States. I say that according to *Fox*, or given what *Fox* says, Obama might have been born outside the US. I say this while I of course know that President Obama was born in Hawaii. The information source in this case contrasts with the above case.

---

closes with some general thoughts on the distinction between first– and third–person epistemic modals.

[5] See Kratzer (2012b/1981).

[6] See Fintel and Heim (2011, §3.3) for the different ways that the ideas of modal bases and conversational backgrounds can be used.

First, the information source is external and non-epistemic. Second, the information source is *static*. What I mean by this is that the *Fox News* website, although it may change, contains static information that isn't updatable in the *epistemic* sense. Third, I don't believe the information source, and in cases like this, I can take a range of attitudes: I may disbelieve, be indifferent, or believe and know it.[7] Such flexibility isn't available in the contrast case. Fourth, the information source isn't perspectival unless we're speaking loosely. Of course we can say things like "From Fox's perspective," but what I think we ultimately mean involves the perspectives of *people* at Fox, not the "perspective" of the content on the website. Because of the non-perspectival, non-epistemic nature of this information source in this situation and similar situations, we'll say that such epistemic sentences are *third person*, or *in* the third person.

The person distinction might bring to mind a notion that Lasersohn (2005) and Stephenson (2007a) have put to work in their accounts of epistemic modals and predicates of personal taste—that of the "autocentric" and "exocentric" positions. These accounts invoke the notion of an index supplemented with a judge parameter. In one way of going about it, a judge parameter is added to the index of evaluation, and an epistemic modal is true (relative to a world and context and possibly more parameters) just in case the judge's beliefs leave the truth of the prejacent open. While it is often the case that the judge is the speaker—the autocentric stance—we can empathize with another's view of the truth of an epistemic modal or predicate of personal taste—the exocentric stance. This distinction can be put to use in various ways, e.g. shifting the judge to the subject of the matrix clause in an attitude report on someone else's epistemic modal or predicate of personal taste claim. Note that this distinction is not the same distinction that I am making. The first- and third-person involve different sources of information, whether of one's own or not, and is not limited to the tastes and information of agents. Put simply and roughly, the autocentric-exocentric distinction distinguishes, respectively, one's own current mental state as the information source of an epistemic modal from another person's current mental state as said source. The first-person--third-person distinction distinguishes, respectively, one's own current mental state as the information source from *either* the mental states of others *or the intentional content of extramental information*—books, magazines, papers, maps, etc.—as said source. That such extramental intentional content can serve as the information source in third-person epistemic modals is perhaps the starkest contrast between the person and centric distinctions. Finally, the autocentric-exocentric distinction is explicitly semantic.[8]

[7] Some consequences of Field's evaluationism will be similar to some of the features I have highlighted here. From an epistemological standpoint, Field's evaluationism divides normative judgments into a norm of evaluation and a belief regarding justification given a norm. On Field's view, our first-person epistemic claims will differ from third-person evaluations regarding some external information and its logical relations to some proposition. I thank an anonymous reviewer for pointing out these connections. See Field (2009, 2018).

[8] We find an analogous situation to the above two, one we're perhaps more familiar with, with arguments. There are at least two different ways we view the premises, the conclusion, and the inferen-

Theorists, often philosophers, discuss epistemic modality as standardly involving knowledge as the information source, as the name "epistemic modal" may suggest.[9] When an agent utters an epistemic sentence that is unmodified (i.e., a BEM), the information source is some kind of knowledge unless there is a third-person information source salient in the context. Thus in standard cases, the information in an epistemic sentence is realistic. In other cases, the information can be third-person and realistic or not. The usual way to involve these third-person information sources is via adverbial modifiers. For example, we can say, "given what the encyclopedia says..." or "given Fox's recent reporting... ."

Importantly for what's to come, we can relate to an epistemic state, frozen in time, whether it be someone else's or a past state of ourselves, in a third-personal way. For example, I may consider what an ancient person knew and consider why, given that knowledge, it was reasonable for her to prefer the geocentric view of the universe prior to the Copernican Revolution.[10] Likewise, someone can consider her past epistemic state and the possibilities it leaves open in a third-personal way. We often heard members of the Bush administration defending their invasion of Iraq on the idea that, given what they knew at the time, Saddam Hussein might have had WMDs, even though, at the utterance time of that epistemic sentence, they knew he didn't.

tial relations between them in an argument. Sometimes the premises are *ours* – we consider them our information and as such we believe and often know them and we draw conclusions from that information. In such a situation, the information that serves as the premises is epistemic, it's dynamic, and it doesn't make sense to say that we're indifferent to or disbelieve it. This situation is analogous to the first-person. But we – and here I mean philosophers – often find ourselves in a different situation: that of mapping logical space. When we consider and give arguments amidst such mapping, we're most interested in the inferential relations between the premises and conclusion, and we're often indifferent toward the truth of the premises and conclusion (or perhaps hostile). In this situation, the information that serves as the premises isn't epistemic, it's *static*, and we're, as noted, often indifferent about the information's veracity. E.g., as we read some paper, we consider what propositions a philosopher's premises or commitments leave open or closed. These premises and commitments aren't *ours* and they're not changeable. We don't "update" them in the same way we update our beliefs. This view is analogous to the second case that involved a third-person epistemic sentence. Unsurprisingly, we find an analogy between arguments and epistemic modals since the standard account of epistemic modality is built on, as Kratzer (2012a/1977, p. 1) says, "a mechanism for drawing conclusions from premises."

[9] As I note in §5.4.2, Egan and Weatherson (2009), in their introduction to a volume on epistemic modality from a mostly philosophical perspective, explicate the three major theories—contextualism, relativism, and expressivism—in terms of knowledge.

[10] This imaginary person couldn't *know* geocentrism is true (since it isn't). Rather, her (at the time) current state of knowledge left open the possibility of geocentrism.

5.2.2 *Conversational Purpose*

Many conversations revolve around locating the actual world in logical space. For example, suppose I'm interested in finding out whether it's raining in New York City (at some particular time), so I check the weather. Divide the set of all possible worlds, i.e. logical space, which "Ω" denotes, into two groups: those where it's raining in New York and those where it's not. Before finding out whether it's raining in New York, I'm unsure where to locate the actual world in logical space between the two groups. Figure 5.1 pictures my epistemic state, and my goal is to move the actual world, which "@" denotes, to the left or right side of the line. On the other hand, we, as conversational participants, may have other interests. For example, we may trace out the consequences of someone's beliefs that we know to be false. For example, suppose we're musing on what the president of the local flat-earth society believes. This person obviously believes that the earth is flat, but we wonder whether his worldview (no pun intended) requires that the moon be flat, too. Let A be the set of worlds where the earth is flat, and let B be the set of worlds where the moon is flat—both metaphysical possibilities. The purpose of this conversation isn't to locate the actual world, which in fact is located outside the circles, but rather to determine (if possible) whether two propositions are compatible according to someone's faulty set of beliefs. Figure 5.2 pictures the state of our musing, deciding between Fig. 5.2a, b.

For epistemic modals, the first-person–third-person distinction allows for different attitudes toward the location of the actual world in logical space. Consult Fig. 5.3 below. The I circle represents the worlds that are compatible with the information, and the P circle represents the worlds at which the prejacent is true. The semantics

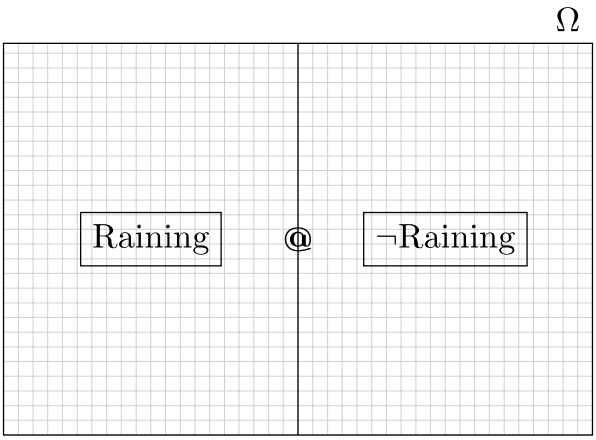

Fig. 5.1 Conversational purpose—locating actual world

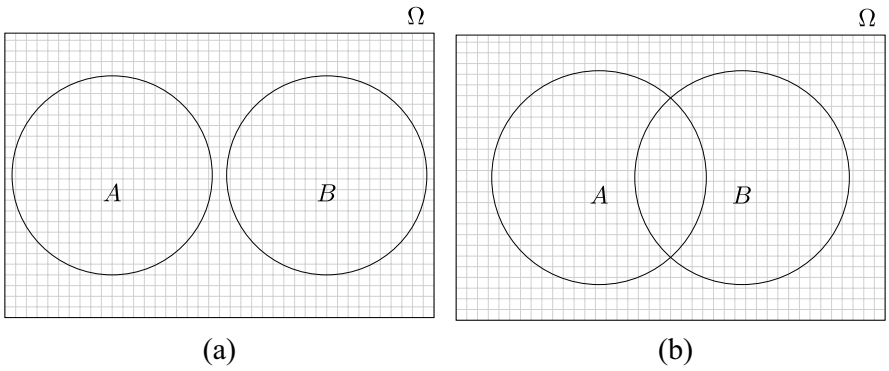

Fig. 5.2 Conversational purpose—filling out logical space. (**a**) Incompatible. (**b**) compatible

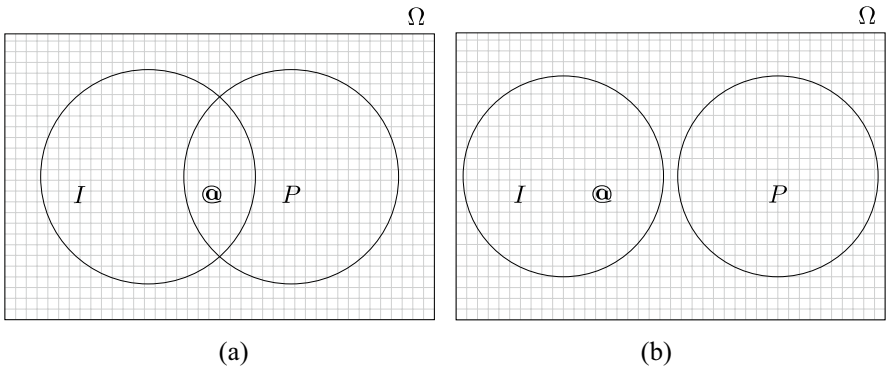

Fig. 5.3 Epistemic modals. (**a**) True first person. (**b**) False first person

for contextualism[11] say that a sentence of the form ⌜It might be that p⌝ is true so long as the information and the prejacent are compatible, i.e. so long as some of the information and prejacent worlds overlap. Erasing the actual world symbols, Fig. 5.3a pictures this semantics. The information source that determines the information involved in the semantics will determine the possible locations for the actual world. Figure 5.3 pictures the options for someone deliberating over the truth of a first-person epistemic sentence. In a first-person epistemic modal, the information is always knowledge, but that knowledge can belong to a single person in a so-called *solipsistic* case or to a group—call that a *group* case. Either way, since knowledge is realistic, the actual world is a member of it. One asserts an epistemic sentence when one is unsure whether to locate the actual world in $I \cap P$, as in Fig. 5.3a or, if the epistemic sentence is false, in which case I and P no longer overlap, in I, as in Fig. 5.3b.

[11] We're ignoring the ordering source.

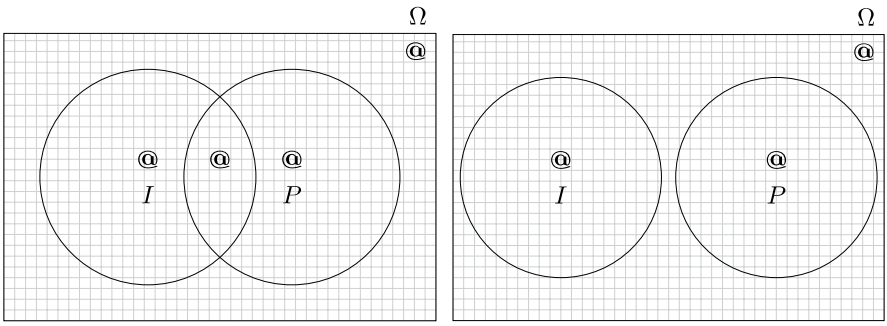

Fig. 5.4 Epistemic modals. (**a**) Compatible third person. (**b**) Incompatible third person

On the other hand, depending on the type of third-person epistemic possibility modal, there are no restrictions on the location of the actual world, as Fig. 5.4 pictures, and Fig. 5.4a, b represent a true and false third-person epistemically modalized proposition, respectively. If the static content is past knowledge (or some other kind of realistic content), then the actual world will be somewhere inside of *I*. But if the content is some other kind of non-realistic content, then the actual world could be anywhere. Saying that the assertor of a third-person epistemic sentence is indifferent to the information's veracity is another way of saying that the speaker is indifferent about the location of the actual world in Ω. Of course, *in fact*, the actual world will be somewhere or other on the diagram. For example, consider the recent United States presidential election in which Joseph Biden defeated Donald Trump, but some sectors of conservative media called into question Biden's win. Suppose I utter the following epistemic sentence a couple of weeks after the election but before the inauguration, yet I know that Biden in fact won: "Given what lots of conservative media is reporting, Biden might have lost." In this context, I'm only interested in the compatibility of the worlds compatible with this media coverage and the worlds at which Biden lost—not in the location of the actual world. Applying Fig. 5.4 to this example, I'm not actively trying to place the actual world in either $I \cap P$ or in $I \backslash P$—I know that the actual world isn't even in *I*. I'm just interested in whether the *I* and the *P* circles overlap.

Before moving on, I want to forestall a possible objection. In what follows, I explain digging in and retracting felicity intuitions in terms of the first-person–third-person and realistic–non-realistic distinctions. One might object that the already established realistic–non-realistic distinction is all that is needed. To foreshadow my response, which I expand on below, one can describe the purpose of a conversation as factual or not. When one attempts to fill out logical space on behalf of the flat earther, there's a sense in which the conversational purpose isn't "realistic"—the information that we're using to fill out what the flat-earther is and isn't committed to is false at the actual world. And while knowledge *is* realistic, there is

a difference between, on the one hand, agents trying to update their first-person beliefs by locating the actual world, and, on the other, filling out logical space based on some information that may or may not be realistic. The realistic–non-realistic distinction does explain the possible locations of the actual world in logical space. But the first–third person distinction captures the difference between information *belonging* to and *being updatable by* an agent, for the first-person case, and the distance an agent has from some information that isn't hers in the third-person case. This idea of what an agent is committed to in a conversation, and possible confusions concerning that commitment, will figure in explaining the difficult cases in §5.3.2. But before getting to those tough cases, let's first clarify the puzzle I am attempting to solve.

5.3 Illustrative Problem Cases

Let's narrow the target of explanation. What we want to explain are dialogues that have the following form:

1. *Disagreement Dialogue Form*
a. (Speaker 1:) An epistemic sentence
b. (Speaker 2:) Disagreement, in some sense, with that epistemic sentence
c. (Speaker 1:) Either (i) or (ii) along with a felicity judgement:
i. Digging in
ii. Folding

In these dialogues, we'll always assume that the disagreer doesn't have the facts wrong.

The number of possible cases explodes quickly depending on how many distinctions we dial in. The previous section canvased some distinctions, and I've alluded to two more. In addition, recall the group–solipsistic distinction and the distinction between whether in (1a) the speaker disagrees with the entire epistemic sentence or just the prejacent. I've also alluded to a spread of cases, some in which (1ci) is appropriate and others in which (1cii) is. Finally, there are some cases in which (1a) is true or false at the time of its utterance by the contextualist's lights.[12] Such a case—in which the initial epistemic sentence is false—are just straightforward cases of disagreement in which a speaker utters a false sentence, a disagreer

[12] For example, suppose, in a logic classroom, I tell a student to assume that $p \rightarrow q$, p are both true. Then I ask the student to consider what well-formed formulas are compatible with this assumption and which are not. The student then says, "Given that $p \rightarrow q$, p are both true, it might be that $\neg q$." This is just a straightforwardly false epistemic sentence. I can either disagree with the entire sentence – saying, "No that's not true – it's not true that given that $p \rightarrow q$, p are both true, it might be that $\neg q$," or just the prejacent: "No that's not true – given our assumptions, q is true."

disagrees, and then the speaker is conversationally obligated to retract. The puzzling cases involve cases in which, in (1a), the prejacent *is* compatible with the information, which either context or an adverbial phrase provides.

To make our task manageable, we need to focus on a limited number of cases. Note that some distinctions don't go together—e.g., third-person epistemic sentences can't be group or solipsistic. Taking that into account, with the following distinctions—first vs. third person, group vs. solipsistic, prejacent vs. modal disagreement, and true vs. false initial epistemic sentence—there are twenty-four combinatorially possible cases, which is a lot of cases. We've just seen that a case in which the initial epistemic sentence is false at the time of utterance is an uninteresting case of disagreement, and aside from some interpretational issues with one of the cases, the group–solipsistic distinction will mostly be irrelevant to what I say. In what follows, we will focus on four main cases.

In the next section, I get the cases on board, say a few things about them, and then zoom out a bit and discuss the cases more broadly. I'll present the cases that require less discussion first.

5.3.1 The Cases

Let's start with a case in which the initial epistemic sentence is third person, the challenger targets the prejacent, and then digging in is infelicitous.

2. *Exploring*:
Context: Magellan and Drake are trying to find Atlantis with a scroll with some vague references to Atlantis' location. Magellan is holding the scroll, and both Magellan and Drake believe the scroll to be accurate. The scroll suggests that Atlantis is in the location that Magellan is pointing.
a. (Magellan, pointing, but not looking:) Given what the scroll says, Atlantis might be right there.
b. (Drake:) Well obviously it's not. We're looking at landless ocean!
c. (Magellan, pointing still:) ? But still, *given what the scroll says*, Atlantis might be right there!

The adverbial phrase makes clear the third person reading, and we've stipulated that the scroll leaves open the prejacent. Drake targets not the inferential relation but rather simply the proposition that Atlantis is in front of them (or something to that effect). So what's wrong with Magellan's digging in? Here, I want to say that, given their joint activity—actively searching for a continent (and not studying, say, maps and scrolls)—Magellan's response is bizarre, perhaps infelicitous. The irrelevance of the digging in to their joint task explains the bizarreness.

The next case is like the last, but the digging in is felicitous.

3. *Bible*:
Context: An atheist Bible scholar is explaining his findings to a lay theist he just met. The findings involve a newly discovered manuscript, which we'll assume leaves open the possibility that Judas was a good man.
a. (Atheist Bible scholar:) Given what this newly discovered manuscript says, Judas might have been a good man.
b. (Theist:) Judas was a treacherous man!
c. (Atheist Bible scholar:) Look *I* have no opinion on Judas. I'm merely saying: *Given what this newly discovered manuscript says*, Judas might have been a good man.

Here, the digging in sounds fine. This is because, as is sometimes the case in scholarly pursuits, the Atheist Bible scholar doesn't really care much about how Judas in fact was. She is only interested in the compatibility of the information from the scroll and Judas being a good man, and about the veracity of the information or the truth of the prejacent, she may be hostile or indifferent. The Theist, on the other hand, is concerned with the actual truth of the prejacent, which explains her response.

Now on to the first-person cases. We'll start with the dialogue type that was involved in the initial case for relativism, and for this reason, I'll call name this dialogue "Usual Relativist." Here we must be careful in dealing with the group–solipsistic parameter. If the initial speaker is included in the group and is mistaken about what possibilities the group's information includes and excludes, then her initial epistemic sentence may be false, so the case would be uninteresting. Thus, we need a case in which the challenger is ignorant with respect to the prejacent. To have a solipsistic case, we need a context in which the participants aren't involved in any joint search, where it's much more plausible that the initial speaker isn't referencing joint knowledge, one in which it would be very strange to prepend "Given what we know, . . ." to the epistemic sentence. Here is such a case:

4. *Usual Relativist:*
Context: Jones is looking for Bob, and Smith is nearby. Jones is the searcher here—it's clear that Jones and Smith aren't jointly searching for Bob, and Smith, though not so much an eavesdropper, is just hanging around.
a. (Smith:) Hey Jones, How you doing?
b. (Jones:) Good, I was just looking for Bob. Oh yeah, I think he might be in his office.
c. (Smith, glances at his office door, on which there's a note:) Oh, no, he's not—there's a note here that says he's left for the day.
d. (Jones:)
i. Huh, I guess I was wrong.
ii. #Sure, but what I said was true.

Finally, onto the case from Fintel and Gillies (2008, p. 81, my numbering) that showed that sometimes it's felicitous to dig in on a first-person epistemic sentence.

5. *vF&G.*
Context: Alex and Billy are looking for keys.
a. (Alex:) The keys might be in the drawer.
b. (Billy, Looks in the drawer, agitated:) They're not. Why did you say that?
c. (Alex:) Look, I didn't say they were in the drawer. I said they might be there—and they might have been. Sheesh.

First, what exactly is Alex saying when she says "[the keys] might have been [in the drawer]"?[13] The intended reading seems to be that the keys' being in the drawer was compatible with the speaker's information at the past utterance time.[14] Alex's digging in, then, is a past-looking, third-person epistemic sentence. The information source is third-person because it's Alex's knowledge frozen at a past time, and Alex claims that her frozen knowledge state at that time left open the key's being in the drawer. Here's another way to put the point. As noted, what we want to say about Alex's response, because of its defensive air, is that she is defending her intelligence. And defending one's intelligence consists in looking *back* on one's knowledge and asserting that it was compatible with some proposition at that time, which, crucially, involves a *different* take on the information. The information source is now static and non-updatable.

Now that we have the cases on board, let's try to say something more general about them.

5.3.2 Zooming Out

Here's the big picture explanation of the felicity intuitions in the cases: conversations can sometimes revolve around locating the actual world in logical space and other times filling out logical space. When there's a tension between the point of a

[13] The answer to this question is somewhat complicated. The epistemic modal operator tends to out-scope the past-tense operator, so we tend to read sentences that involve the two operators as referencing the *current* information about a *past* possibility. While "might have been" is often read as non-epistemic, an epistemic interpretation of "might have been" appear in sentences such as "Jones isn't a philosopher, but he might have been" in which I know Jones is not a philosopher, but the *current* information leaves open whether—at some time before now—he was. The corresponding interpretation isn't what we're looking for in *vF&G*. This would mean that Alex, in (5c), is asserting that the keys aren't in the drawer, but the current information leaves open whether, at some time before the assertion time, the keys were in the drawer. However, the context makes clear that the keys weren't in the drawer at some relevant past time—maybe they were days ago, but an assertion with that meaning would be irrelevant.

[14] The information involved in Alex's initial epistemic sentence, given the facts of the case, can either be Alex's knowledge alone or both Alex and Billy's. It's clear that Alex's knowledge is compatible with the prejacent when she makes her initial assertion, and Billy looked for the keys *after* Alex's initial assertion, so Billy didn't know where they were at the time of the utterance. So both group and solipsistic readings are compatible with the case. My explanation is compatible with either option. I do, however, think it's much more plausible that this is a solipsistic case. As I note shortly, Alex's digging in involves defending her intelligence, and I think the idea of defending one's intelligence is more plausible on a solipsistic rather than a group reading.

participant's assertion and the point of the conversation, the participant's assertion will be irrelevant, sometimes so irrelevant that that defectiveness explains the infelicity.

The first-person naturally pairs with interest in locating the actual world in logical space. Cases like *Usual Relativist* are paradigmatic here. As noted, when I say, "Jones might be in his office" amidst searching for Jones, I'm trying to locate the actual world in logical space. When I say that sentence, what I know leaves open my world being a Jones-in-office world. But given my activity—searching for Jones—I would like further to narrow down Jones' location—i.e., further to narrow down the location of the actual world, finding out whether or not Jones is in his office. Applying Fig. 5.3a to this example, where P is the proposition that Jones is in his office (at some time), I'm not sure which side of the line to place the actual world, but I'm actively trying to find out.

Just as epistemic sentences pair with a certain attitude toward the location of the actual world, so can conversations: some revolve around locating the actual world in logical space; others, around filling it out. But when we mismatch the conversational attitude with the wrong kind of epistemic sentence, tension results because the epistemic sentence isn't relevant to the conversation. We can now say what explains the felicity intuitions: We get infelicity if the tension is great enough. "Tension" here means a difference between the point of an assertion and a conversation. Note that when I say "tension," I'm not introducing new or interesting philosophical terms here to do major work. I'm merely saying what irrelevance is in other words: Conversations tend to revolve around a given topic, and infelicity ensues if someone says something off-topic enough. The point I'm making is to say first that irrelevance explains the typical infelicity of digging in, and second, to point out the source of the irrelevance. To put the second point differently: To say why a given instance of digging in is irrelevant enough to result in infelicity, I need to say why such an assertion is off-topic, and here's why: The topic is locating the actual world; digging in isn't on that topic.

Let's see how this works out with our cases. First, let's explain the different felicity data in the *Usual Relativist* cases versus the *vF&G* case. In the *vF&G* case, the context clarifies that the participants are actual-world hunting. In other words, they are trying to figure out which side of the line to place the actual world on Fig. 5.3a, which is how logical space looks at the time of the initial utterance. After Billy's challenge, the actual world moves to the left, in $I \backslash P$. But the challenger introduces a special feature of this context: her assertion is somewhat heated and accusatory. This prompts Alex's response, which in terms of our diagram, says that, although she *currently* knows that logical space is such that $@ \in I \backslash P$, it was reasonable for her to place the actual world on the line when she uttered her initial epistemic sentence. Alex's assertion in her digging in is in the third person: She's defending the reasonableness of I and P overlap. Being third person, Alex's digging in also has the effect of abandoning the hunt for the actual world. I said that we get infelicity if the tension is great enough. Here, the point of the conversation is initially to find keys, and Alex's digging in doesn't contribute to that. But there is less tension because of the special context—there is a reason for Alex's response, viz. Billy's need to relax a bit.

I've given some context to *Usual Relativist*, but it's no accident that similar cases are glossed with little to no context. Like *vF&G*, *Usual Relativist*, starts in a context where the participants attempt to locate the actual world. But unlike *vF&G*, the challenger doesn't take an accusatory tone. Because of this, Jones has no reason to defend his intelligence—doing so would be bizarre. In other words, asserting the third-person epistemic sentence implicated in digging in would introduce an intolerable tension between the point of the conversation (locating the actual world—i.e. finding Bob) and the point of the epistemic sentence (not locating the actual world). In yet other words, digging in would be so irrelevant as to be infelicitous.

We have two third-person cases to explain: *Bible* and *Explorer*. *Bible* is quick to explain: The Atheist Bible Scholar only cares about the overlap between *I* and *P*, not in locating the actual world, which the Theist is interested in, as her reply shows. The Atheist Bible Scholar's digging in is a way of stressing that she's only interested in *I–P* overlap, not in locating the actual world. If there is any tension in this conversation, it's that the Theist is somewhat confused about the point of the Atheist Bible Scholar's assertion, but the confusion is mild, and the Atheist Bible Scholar simply lets the Theist know what she is up to. Another way to put this is that there's no clear point, and some confusion, over what the conversation revolves around, and the digging in resolves the confusion.

Explorer is interesting in that it somewhat mimics *Usual Relativist*. There are two reasons for this: The context is one in which the conversational participants attempt to locate the actual world, and the conversational participants *take* the information to be realistic. Drake and Magellan presume the map to be veracious, and the feel of the context is that it's somewhat high stakes—they're on the high seas searching for a lost continent. If Magellan were to dig in, he would abandon his search for the actual world to consider how logical space looked at some past time. The result would be a tension between a conversational context that revolves around locating the actual world and a sentence that would abandon that purpose—a sentence whose point would be to consider the content of a map they currently know is faulty. Such tension would be too great, and so the digging in would be irrelevant to the point of infelicity.

In summary, here are the diagnoses of the cases:

Infelicitous

- Both *Usual Relativist* cases and *Exploring*. Context involves locating actual world. Digging in involves consideration of logical space at some past time with no special prompt to do so.

Felicitous

- *Bible*. Context is somewhat unclear. Digging in resolves the unclarity.
    - *vF&G*. Context involves locating actual world. Digging in involves consideration of logical space at some past time, but there is a special prompt to do so.

5.4 Closing Remarks

5.4.1 Areas for Further Inquiry

Knobe and Yalcin (2014) and Khoo (2015) present empirical work supporting contextualism over relativism, to which Beddor and Egan (2018) responds from a relativist perspective. Beddor and Egan introduce the view they call "flexible relativism", which they describe as follows:

> The proposition that's expressed by an utterance of Might p is a centered worlds proposition—the one corresponding to the property: lacking information that rules out p. Asked to deliver a verdict about the truth or falsity of the utterance, there are different questions one could be seeking to answer. One could be offering an answer to the question, is the proposition expressed true of me (the assessor)? Or one could be offering an answer to the question, is the proposition expressed true of the speaker? (Beddor & Egan, 2018, p. 9)

Assessors will gravitate toward different contexts of assessment depending on which question is currently under consideration. In light of this, Beddor and Egan (2018, p. 10) pose the following pragmatic constraint, the Question Under Discussion (QUD) constraint:

> Suppose someone in a conversational context c is assessing an utterance of a BEP for truth or falsity. Ceteris paribus, they will be inclined to assess the utterance using whichever context of assessment is most relevant to answering the QUD in c.

My explanation of *vF&G* involved considering a past epistemic state and the possibilities it left open, similar to Beddor and Egan's cases involving a question under discussion of an agent's competence. In their cases, the prejacents of the epistemically modalized propositions are false, and they elicit various judgments concerning the competence of those that utter these propositions. Beddor and Egan may subsume the conversational goals I consider as another QUD, but the person distinction isn't yet another QUD. And as Beddor and Egan note (fn.32), a natural extension of their work will turn to embedded contexts, which I have considered here. But unlike them, I have not put my judgments to an empirical test, which would be a possible next step.

5.4.2 Reflections on the Person Distinction

A modality is alethic if, but only if, necessarily, the actual world is a member of the base. If *must-p* or *might-p* are true, and the modality is alethic, then, necessarily, the actual world will be a p world. And the following inferences are characteristically valid for alethic modalities:

$$p \vDash @\, p; @\, p \vDash \cdot\, p.$$

The focus on realistic information has, I think, made us forget that epistemic modals are sometimes alethic and sometimes not. It's hard to overstate how many critical papers in at least the philosophy literature suggest that knowledge is the exclusive base for epistemic modality. Egan and Weatherson (2009), in their introduction to the only philosophical anthology on epistemic modality, define, e.g., contextualism in terms of knowledge and expound on considerations that only make sense in terms of knowledge—e.g., the so-called "speaker inclusion constraint," which requires that the speaker's knowledge is a part (proper or improper) of the knowledge that determines the base. Such a consideration only makes sense, obviously, if something epistemic determines the base.

If the information source is realistic, then the modal will be alethic; otherwise, not. Epistemic modality houses both alethic and non-alethic versions of itself. And the alethic–non-alethic distinction gives us another angle on some of what we've explained. When a modality is non-alethic, it comes as no surprise that one can dig in in the face of a prejacent challenge. For example, if I point to a crosswalk on a new road and say that pedestrians can cross there, and someone informs me that it's not true that pedestrians cross there (since none have), then I can felicitously reiterate my deontic sentence. The same sometimes goes for an epistemic sentence that is initially asserted in the third person.

To close, let us briefly consider the nature of the first–third-person distinction. I think we best understand the distinction as non-semantic, and I've appealed to non-semantic notions, viz. relevance, in my main argument. I've also tried to speak in terms of different information *sources* in distinguishing first- and third-person epistemic sentences rather than different information *simpliciter*. We should understand the difference between first- and third-person epistemic sentences as a difference in the way information is determined rather than in terms of different *kinds* of content.

The kind of case that convinces me is one in which agents commit differently to the same content, assuming the simple-minded version of contextualism that I've been assuming in this paper. The difference in commitment arises from different information sources, either first or third person. In other words, we can construct a case in which an agent's knowledge and some external information source determine the same information involved in an epistemic sentence, yet the commitment facts differ. Here is such a case.

Suppose I come into existence, *tabula rasa*, and read the entire *Encyclopedia Britannica*, which exhausts the entirety of my knowledge and belief. Suppose that *Britannica* strongly suggests that Shakespeare wrote all of the plays attributed to him but also suggests there is a modicum of evidence that Francis Bacon wrote a few of the plays. Now suppose I assert the following:

6. (I:) Bacon might have written *Romeo and Juliet*.

My information leaves open whether Bacon wrote Romeo and Juliet. Suppose that you have also read *Britannica*, but it doesn't exhaust your knowledge and belief.

You are in fact a Shakespeare scholar and have considered all the information, and you have ruled out that anyone but Shakespeare wrote *Romeo and Juliet*. You say:

7. (You:) Given what *Britannica* says, Bacon might have written Romeo and Juliet.

Here the same intensional content, *Britannica*, determines the information that (it is asserted) leaves open the possibility that Bacon wrote *Romeo and Juliet*. For me, however, that content exhausts my epistemic state. For you, it is the content of a set of encyclopedias. According to a very simple version of contextualism, (6) and (7) determine the same propositions, yet these different assertions come with different commitments. In making your assertion, you play the role of arbiter. You're considering the inferential relation between some static information and a proposition. This being the case, you can follow up your claim with, "but of course Shakespeare wrote *Romeo and Juliet*." On the other hand, I, issuing a first-person epistemic sentence, cannot give the same follow-up to my assertion.

Finally, let's put this point in general way. We've repeatedly made the point that we can reference particular information sources that determine the information involved in the proposition an epistemic sentence expresses. And though I haven't spoken in terms of the formal semantics of epistemic modality much up to now, in closing, I bring it in because I think it makes this point especially clear. According to standard possible worlds semantics (leaving out the ordering source), we associate the information with a set of possible worlds compatible with that information, and the possibility modal particles *might* or *may* semantically contribute an existential quantifier over that set of worlds. We state the truth conditions in terms of existential quantification over the information worlds: An epistemic sentence is true just in case there is at least one world in the information worlds at which the prejacent is true. If we identify these truth conditions with the proposition an epistemic sentence expresses, then, at the level of the proposition, different information sources can determine the same set of worlds, as in the above case. And we can't, by inspecting this or that proposition that an epistemic sentence expresses, determine how a speaker or hearer may relate to that proposition vis-à-vis how they relate to the actual world. There is, so to speak, no road backward from the proposition an epistemic sentence expresses to the information source involved in that proposition. I'm unsure what to make of the equivalent point, *mutatis mutandis*, for other flavors of modality—that, e.g., different rule sources can determine the same quantificational domain for deontic modals. But for epistemic modality, it's significant.

Acknowledgments Thanks to Phillip Bricker and Seth Cable for their helpful comments.

References

Beddor, B., & Egan, A. (2018). Might do better: Flexible relativism and the QUD. *Semantics and Pragmatics, 11*.

Cappelen, H., & Hawthorne, J. (2009). *Relativism and monadic truth*. Oxford University Press.

DeRose, K. (1991). Epistemic possibilities. *The Philosophical Review, 100*(4), 581–605.

Dorr, C., & Hawthorne, J. (2013). Embedding epistemic modals. *Mind, 122*(488), 867–914.
Dowell, J. L. (2011). A flexible contextualist account of epistemic modals. *Philosophers' Imprint, 11*(14).
Egan, A. (2007). Epistemic modals, relativism and assertion. *Philosophical Studies, 133*(1), 1–22.
Egan, A., & Weatherson, B. (Eds.). (2009). *Epistemic modality*. Oxford University Press.
Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 131–170). Oxford University Press.
Field, H. (2009). Epistemology without metaphysics. *Philosophical Studies, 143*(2), 249–290.
Filed, H. (2018). Epistemology from an evaluativist perspective. *Philosopher's Imprint, 18*(12).
Hawthorne, J. (2007). Eavesdroppers and epistemic modals. *Philosophical Issues, 17*(1), 92–101.
Khoo, J. (2015). Modal disagreements. *Inquiry, 58*(5), 511–534.
Kneer, M. (2022). Contextualism vs. relativism: More empirical data. In J. Wyatt, J. Zakkou, & D. Zeman (Eds.), *Perspectives on taste: Aesthetics, language, metaphysics, and experimental philosophy* (pp. 109–139). Routledge.
Knobe, J., & Yalcin, S. (2014). Epistemic modals and context: Experimental data. *Semantics and Pragmatics, 7*(10), 1–21.
Kratzer, A. (2012a). What 'must' and 'can' must and can mean. In A. Kratzer (Ed.), *Modals and conditionals: New and revised perspectives* (pp. 1–20). Oxford University Press.
Kratzer, A. (2012b). The notional category of modality. In A. Kratzer (Ed.), *Modals and conditionals: New and revised perspectives* (pp. 21–69). Oxford University Press.
Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy, 28*(6), 643–686.
von Fintel, K., & Gillies, A. S. (2008). CIA leaks. *Philosophical Review, 117*(1), 77–98.
von Fintel, K., & Gillies, A. S. (2011). Might made right. In B. Weatherson & A. Egan (Eds.), *Epistemic modality* (pp. 108–130). Oxford University Press.
von Fintel, K., & Heim, I. (2011). *Intensional semantics* (Spring 2011 Edition). Unpublished manuscript.
MacFarlane, J. (2003). *Epistemic modalities and relative truth*. Unpublished manuscript.
MacFarlane, J. (2009). Epistemic modals are assessment-sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 145–178). Oxford University Press.
Marques, T. (2018). Retractions. *Synthese, 195*(8), 3335–3359.
Schaffer, J. (2011). Perspective in taste predicates and epistemic modals. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 179–226). Oxford University Press.
Stephenson, T. (2007a). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy, 30*(4), 487–525.
Stephenson, T. (2007b). *Towards a Theory of Subjective Meaning*. Doctoral dissertation, Massachusetts Institute of Technology.