

Knowledge by Constraint*

Ben Holguín

October 2021

Abstract

This paper considers some puzzling knowledge ascriptions and argues that they present *prima facie* counterexamples to credence, belief, and justification conditions on knowledge, as well as to many of the standard meta-semantic assumptions about the context-sensitivity of ‘know’. It argues that these ascriptions provide new evidence in favor of contextualist theories of knowledge—in particular those that take the interpretation of ‘know’ to be sensitive to the mechanisms of constraint.

1 Introduction

Start with a case:

History Exam Two students, Peggy and Pete, have purchased and subsequently memorized answer sheets from their corrupt teaching assistant, Roger. However, moments before the exam Roger shares an unfortunate discovery: one of Peggy or Pete was given answers to the wrong exam, but he doesn’t know who it was. All Roger knows is that the two answer sheets give different answers to every question.

...

The first question of the exam reads: “In what year did the Berlin Wall fall?”. Peggy’s answer sheet says b: 1989, Pete’s says c: 1991. Peggy is thus the one with the good answers.

Putting ourselves in Peggy’s shoes, here are some natural sounding things we could think or say:

- (1) If my answer sheet is good, then I know what the answer to this question is.
- (2) One of us knows what the answer to this question is. (I hope it’s me.)

*Forthcoming in *Philosophical Perspectives*. Thanks to Bob Beddor, Sam Carter, Mala Chatterjee, Cian Dorr, Jane Friedman, Jeremy Goodman, Dan Hoek, Paul Horwich, Simon Goldstein, Nico Kirk-Giannini, Harvey Lederman, Andrew Lee, Sam Lee, Matt Mandelkern, Jake Nebel, Jim Pryor, Stephen Schiffer, Ginger Schultheis, and Jake Zuehl for helpful discussion and comments. Special thanks to Kyle Blumberg and Trevor Teitel for continued feedback over the course of the project.

- (3) Either Pete has the good answers, or I know what the answer to this question is.
- (4) I might know what the answer to this question is; it depends on whether I got the good answers or not.

And putting ourselves in Roger (the TA)'s shoes, it seems we could think or say:

- (5) Whoever has the answer sheet that says '1: b' knows the answer to that question.

Upon discovering that it is Peggy who has the good answers, we could then say:

- (6) Peggy knows what the answer to the first question is.

But now we have a puzzle. Taking the judgments about (1)–(6) at face value, it looks to follow from the fact that Peggy has the good answers that she knows what the answer to the first question is (likewise for every other question). This kind of knowledge, if genuine, is rather strange. Among other things, it appears to be in tension with the following widely accepted principles about knowing:

CREDENCE If S knows that p, then S's credence that p is greater than .5.

BELIEF If S knows that p, then S believes that p.

JUSTIFICATION If S knows that p, then S is in a position to justifiably believe that p.

Start with **CREDENCE**. Intuitively, if Peggy is rational then her credence in the proposition that the answer to the first question is b can be no higher than .5. She knows that if and only if she got the good answers is b the answer to the first question, and she knows that the odds she got the good answers are no higher than .5. Granted, Peggy knows that Pete's answer sheet is equally likely to say a, c, or d, and for that reason may be rationally compelled to put b rather than one of the other answers. But on the question of 'Is it b or not-b?', she's essentially facing a coin flip.

The following much weaker principle can be rejected along similar lines:

MINIMAL CREDENCE If S knows that p, then S's credence that p is greater than .001.

And that is because we can reimagine *History Exam* to make the odds more as less as low as we want without affecting the relevant judgments. Just imagine, for instance, that there are 1,000 possible answers to a given multiple choice question and that there are 1,000 different students taking the test, each having purchased a different answer sheet from Roger, and that one and only one of the students has the good answers (and that Peggy knows all this). Supposing Peggy is the student with the good answers, (1)–(6) remain acceptable, even though now Peggy is aware that the odds that the answer to the first question is b are no higher than .001.

Now for **BELIEF**. If we think that principles like **CREDENCE** (or at least **MINIMAL CREDENCE**) hold for belief, then the above considerations provide a direct argument against **BELIEF**: Peggy

knows that the answer to the first question is *b*, but her credence in that proposition is not greater than .5 (and could be as low as .001). Belief implies credence greater than .5 (or .001), so knowledge cannot require belief. It is also worth noting that with the exception of the fact that Peggy intends to fill in ‘*b*’ for her answer, Peggy lacks any of the standard hallmarks of belief. She would not assert that the answer to the first question is *b*, nor would she bet on more than even odds on it, nor would she be surprised to discover that the answer is something other than *b*. If Peggy believes the answer to the first question is *b*, it is presumably in a sense that has not been on the minds of epistemologists who have defended principles like BELIEF.¹

Next, JUSTIFICATION. If Peggy were to believe that the answer to the first question is *b*, would she be justified in doing so? For roughly the reason sketched in the discussion of CREDENCE, the answer seems like an obvious no. The entirety of Peggy’s evidence for the proposition that the answer to the first question is *b*—at least on any intuitive way of understanding ‘evidence’—is that her answer sheet says as much. But Peggy also seems to know that her answer sheet has at most a .5 chance of being correct. And remember: the case can be modified to make that number as small as one wants without affecting the relevant linguistic judgments.

In short, the uses of ‘know’ raised to salience by *History Exam* cry out for explanation. If the intuitions about (1)–(6) are taken at face value, then we have a case that presents a strong challenge to a number of widely accepted epistemological principles. Moreover, supposing the principles are ultimately false, the case seems not to suggest any straightforward explanation of *why* they would have seemed so plausible in the first place. Peggy’s knowledge, if genuine, is radically unsupported by what we’d normally call ‘the evidence’, and is in most relevant respects detached from anything philosophers normally call ‘belief’.

The fact that we can think and talk in these ways has not been given much attention in the epistemological literature.² This paper aims to fill that lacuna. It will work under the assumption that our intuitive judgments about sentences like (1)–(6) ought to be taken at face value, with the goal of exploring the semantic, meta-semantic, and epistemological consequences of adopting a theory of knowledge that is capable of accommodating them. Though §2 (and to some degree §4) will defend this assumption at some length, I won’t purport to give a knock-down argument in its favor. Those who remain skeptical that the judgments ought to be taken at face value should treat the paper as a test of assumptions, seeing how many of the standard views on knowledge must bend or break to account for these ways of thinking and talking.

With that, the paper is structured as follows. I start by arguing that the judgments brought out by *History Exam* are not susceptible to simplistic, non-semantic kinds of explanations (§2). I then turn to developing a contextualist theory of knowledge that can account for the data (§3).

¹ As we will see in §5, there is reason to think that this may be precisely the sense in which BELIEF is true.

² The closest comparison I know of are some cases of “loose knowledge” discussed by Hawthorne (2000, pp. 202-3; 2004, pp. 68-9). However, Hawthorne’s cases do not involve subjects who are aware of their impoverished epistemic circumstances—as Peggy is in *History Exam*—nor do they involve judgments about knowledge ascriptions embedded under conditionals, quantifiers, disjunctions, or modals—as in (1)–(4) respectively. As will be seen in §§3–4, these features of *History Exam* and related cases are of special theoretical interest.

Issues concerning its meta-semantics are discussed at length (§4). After that I examine some patterns in the data that the view has trouble explaining (§5), and defend an amendment that accounts for them—one that revives a version of the BELIEF condition on knowledge. (§6) concludes with a discussion of what to make of principles like CRE-DENCE, BELIEF, and JUSTIFICATION in light of the paper’s arguments.

2 Objections to taking the judgments at face value

This section will discuss four objections to the claim that our intuitive judgments about *History Exam* ought to be taken seriously for the purposes of semantic and epistemological theorizing. In doing so it will expand on the patterns underlying the knowledge ascriptions of interest.

2.1 First objection: the generality of the phenomenon

A common first blush reaction to *History Exam* is to chalk up the strangeness of (1)–(6) to the distorting effects “exam cases” have on our judgments about knowledge. The worry is that exam settings tempt speakers to treat ‘know’ as meaning something along the lines of ‘can correctly answer’. And since there is no reason to expect ‘can correctly answer’ to abide by any of CRE-DENCE, BELIEF, or JUSTIFICATION, we shouldn’t be so surprised if, when thinking about *History Exam*, we talked about knowledge as if it didn’t either.

I will argue against “reinterpretation” strategies along these lines in more detail when I turn to the fourth objection (§2.4). For now, I simply want to point out the phenomenon raised to salience by *History Exam* is a fairly general one, and has nothing special to do with exam settings. Here is a case in illustration:

Office Meeting Don and Betty have both been asked to attend a company meeting by their boss, Roger. However, shortly before the meeting Roger lets them know that one of them was given the wrong room number for the meeting, that he doesn’t know who it was, but that if either of them gets to the room and see no one is there, he or she should give him a call. (As it turns out, Don was correctly told room 101, Betty incorrectly room 110.)

I submit that it would be natural for Don to say or think things like:

- (1_R) If I got the right room number, then I know where to go for the meeting.
- (2_R) One of us [Betty or Don] knows where to go for the meeting.
- (3_R) Either Betty knows where to go for the meeting, or I do.
- (4_R) I might know where to go for the meeting; depends on which person Roger gave the wrong room number.

And since *in fact* Don got the right room number, it follows that he knows where to go for the meeting. And once Roger finds out that he screwed up with Betty, not Don, he can think and say things like:

(6_R) Don is the one who knows where to go for the meeting.

The analogies to *History Exam* should be obvious.

Next, if we follow Williamson (2000, ch. 1) in taking factive mental stative operators like ‘remembers that’, ‘sees that’, etc., to entail their ‘knows that’ counterparts, then cases like the following present yet more instances of this kind of knowledge:³

Memory Experiment Joan and Megan are participating in a trial of a drug whose primary effect is to swamp its subjects with an extraordinary number of fake “memories” of the events of the past 24 hours. One of the subjects will get the drug, while the other will get a placebo. Who gets which is determined by a coin-flip whose result is known only to the experimenters.

During the experiment Joan and Megan are both (separately) asked ‘Do you remember what you ate for dinner yesterday?’ Joan appears to remember that she ate fish; Megan appears to remember that she ate spaghetti. As a matter of fact it was Joan who got the placebo and Megan who got the drug.

I submit it would be natural for Joan to say or think things like:

- (1_M) If I got the placebo, then I remember what I ate for dinner last night.
- (2_M) One of us remembers what she ate for dinner last night.
- (3_M) Either Megan got the placebo, or I remember what I ate for dinner last night.
- (4_M) I might remember what I ate for dinner last night; it depends on whether I got the placebo or the drug.

Since Joan *in fact* got the placebo, it follows that she remembers (and thus knows) what she ate for dinner last night. And since the experimenters know who got the placebo, they can say things like:

(6_M) Joan remembers what she ate for dinner last night.

Examples can be multiplied with ease. As such, the underlying phenomenon brought out by *History Exam* should not be thought to have any essential connections to the idiosyncrasies of our judgments about knowledge in exam settings. Something more general is going on.

³ Blumberg & Holguín (2019) discuss at length similar cases involving emotive factive verbs like ‘surprise’ and ‘regret’.

2.2 Second objection: wide-scoping

A popular reaction to the *conditional* knowledge ascriptions—e.g., (1) ('If my answer sheet is good, then I know what the answer to the first question is')—is to claim that our intuitions about these sentences rest on non-obvious subtleties concerning the scope of 'know' at logical form. In particular, the thought is that the natural readings of (1) are *wide-scope* readings—that is, readings on which the attitude verb takes scope over the whole conditional, rather than just the embedded -wh clause. According to this view, the true proposition we associate with (1) is better represented by:

(1_w) I know whether: if my answers are good, then the answer to the first question is b.

And it is obvious that it does not follow from (1_w) plus the fact that Peggy's answers are good that Peggy knows what the answer to the first question is.

However, there are at least three reasons to think that as a general account of the phenomenon, wide-scoping is a non-starter.

First, the simplest one: there is no reason to expect scope ambiguities to explain the full range of data. *History Exam* raises to salience a number of sentences that are clearly *not* indicative conditionals that seem to pose exactly the same problem (1) poses on its narrow-scope interpretation: (2) is a quantified knowledge ascription, (3) is a disjunctive knowledge ascription, (4) is a knowledge ascriptions embedded under a modal, and (5)–(6) are plain old unembedded knowledge ascriptions. So even supposing we could in principle explain away the true readings of (1) in terms of something like a scope ambiguity, it is not at all clear how the story is supposed to be extended to any of the rest of the data.

Second, the wide-scope response founders on third-person analogs of (1). Consider a slightly altered version of *History Exam* in which Peggy sells her answers to a third student, Ken, who believes falsely that he, Peggy, and Pete are all studying from the same answer sheet. Taking Peggy's perspective in this new setting, (1+) is fine even though its wide-scope analog (1+_w) seems false:

(1+) ✓ If Pete got bad answers, then Ken and I both know what the answer to the first question is.

(1+_w) ? Ken and I both know whether: if Pete got bad answers, then the answer to the first question is b.

(1+) is puzzling in the same ways (1) is, yet the former evidently cannot be explained away in terms of a scope ambiguity.

Third and finally, consider a conditional ascription like:

(7) ✓ If Pete doesn't know what the answer to the first question is, then I do.

(7) seems to me as felicitous as any of (1)–(6). Indeed, it’s basically just another way of saying what (2) says (‘One of us knows what the answer to the first question is’). But assuming ‘know’ only has its standard readings, the wide-scope analog of (7) is clearly false:

(7_w) ? I know that: if Pete doesn’t know what the answer to the first question is, then the answer to the first question is b.

If Pete has the good answers, then the answer to the first question can’t be b. But supposing that ‘know’ never takes the strange readings I’m claiming it sometimes takes, it would follow that regardless of whether Pete has the good answers, ‘Pete doesn’t know what the answer to the first question is’ expresses a truth. And since (by hypothesis) Peggy doesn’t know whether the answer to the first question is b, it would follow that for all Peggy knows, ‘If Pete doesn’t know what the answer to the first question is, then the answer to the first question is b’ expresses a falsehood. That is: Peggy would know that the conditional’s antecedent is true, but not whether its consequent is. And so (7_w) should *not* seem like the kind of thing Peggy could think or say. But it is a perfectly fine thing for her to think or say. So wide-scoping does not explain the good readings of (7).

2.3 Third objection: knowledge-wh vs. knowledge-that

As some readers may have noticed, each of *History Exam*’s (1)–(6) involves ‘knowledge-wh’ rather than ‘knowledge-that’. This is because the knowledge-that counterparts of (e.g.) (1), (3), (4), and (6) tend to sound less natural:⁴

(1*) ? If my answer sheet is good, then I know that the answer to the first question is b.

(3*) ? Either Pete got the good answers, or I know that the answer to the first question is b.

(4*) ? I might know that the answer to the first question is b; it depends on whether I got the good answers or not.

(6*) ? Peggy knows that the answer to the first question is b.

But the principles for which I claim these data raise problems—CREDCENCE, BELIEF, and JUSTIFICATION—are all stated in terms of knowledge-that. So perhaps what cases like *History Exam* show is not that one can know without believing or having justification to believe, but rather that it is possible to know *what* the answer to a multiple-choice question is without knowing of the correct choice that *that* is the answer to the question. More generally, these cases might be taken to show that -WH TO -THAT is false (here ‘Q’[?] ranges over -wh questions):

⁴ (2) and (5) are omitted because neither admits of a straightforward ‘know-that’ translation. I should note that not everyone seems to have the judgment that the ‘know-that’ ascriptions are worse than the ‘know-wh’ ascriptions, but many do. I myself find them a bit worse, with the exception of (1*), which I find perfectly natural.

-WH TO -THAT: In every context c : If ‘S knows $Q^?$ ’ is true in c , then there is a contextually salient answer to $Q^?$, p , such that ‘S knows that p ’ is true in c too.

Were it true, the claim that *History Exam* presents a counterexample to -WH TO -THAT would be interesting in its own right.⁵ But there are at least two reasons to think that that’s not what is going on here.

First, if -WH TO -THAT were false, then we should expect (8) to have a true uniform interpretation:

- (8) ?? Although Peggy knows what the answer to the first question is [namely b], she does not know that the answer to the first question is b .

But instead it just sounds like a contradiction. It is not at all obvious why this would be so were *History Exam* a simple counterexample to -WH TO THAT.

Second, it really isn’t so hard to hear the true readings of the ‘know-that’ ascriptions. Three quick points in illustration:

- A) The ‘knows-that’ analogs of (1)–(6) read perfectly naturally when emphasis is placed on any of ‘Peggy’, ‘the first question’, or ‘ b ’.

- B) Once we’ve anchored ourselves into a context where the ‘knows-wh’ ascriptions sound fine, the ‘knows-that’ ones start to sound fine too. For example, the first response to (9) is entirely appropriate, while the second seems bizarre:

(9) Q: Does Peggy know what the answer to the first question is?

A₁: ✓ Yes, she knows that the answer to that question is b .

A₂: ?? Yes, but she does not know that the answer to that question is b .

- C) Finally, the natural answer to a question like ‘Which of the students knows that the answer to the first question is b ?’ is ‘Peggy’, not ‘Neither of them’. And against thinking that the preference for ‘Peggy’ is forced because the question *presupposes* that at least one of the students knows that the answer to the first question is b , notice that the natural answer to questions like ‘Which of the students is certain the answer to the first question is b ?’ is ‘Neither of them’, not ‘Peggy’.

One who holds the view that the ‘know-that’ analogs of (1)–(6) are all false must therefore do more than just explain away the appeal of -WH TO -THAT. One must also explain why it is so easy to get into a frame of mind on which those ascriptions are intuitively true. It is not clear what sort of story can be told here if in fact the ‘know-that’ ascriptions are context-invariantly false.

⁵ With the exception of Farkas (2016, 2017), I know of no theorist who denies -WH TO -THAT. Indeed Farkas herself claims that she is the only person in the literature who denies it (2016, p. 857). Note that her putative counterexamples to -WH TO -THAT are not at all like *History Exam*.

2.4 Fourth objection: loose speech

The last objection I will consider claims that strange ascriptions like (1)–(6) are mere *loose speech*.⁶ That is to say: as uttered in context, none of (1)–(6) is literally true. Impressions to the contrary are due to the fact that (i) each can be used to communicate a true proposition other than the one determined by the literal meaning of the sentence in context, and (ii) our intuitions are primarily tracking these (non-literal) communicated propositions. What is the true but non-literal proposition communicated by, e.g., (6, ‘Peggy knows what the answer to the first question is’)? Well, that depends on one’s preferred version of the objection. Perhaps it is that Peggy knows for the purposes of the exam what the answer is. Or perhaps it is that Peggy will behave as one who knows what the answer is. Or perhaps it is something else entirely. I think there are specific grounds on which one could quibble with any proposed paraphrase of this sort.⁷ But by way of making a more substantive case against the loose speech response, it is worth comparing its present application to an application that is less controversial: namely, accounting for so-called non-factive uses of ‘know’.

As is well known, there are certain contexts in which utterances of ‘S knows that p’ can seem basically fine even when it is common knowledge that p is false. Here are two examples, the first from Hazlett (2010, p. 501):

- (10) Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.
- (11) Stranded under enemy fire, I knew I was going to die there. But then the air support came in and I somehow made it through.

Should we take the fact that these sentences sound basically fine to show that our best theory of knowledge should reject the entailment from ‘S knows that p’ to p? I follow orthodoxy in saying no.⁸ I think these are just instances of loose speech. And the reason I am comfortable dismissing *these* felicitous uses of ‘know’ as loose speech, while not (1)–(6), is that in the non-factive case, the loose speech diagnosis can be motivated on independent and principled grounds.

⁶ I use the label ‘loose speech’ somewhat loosely, and intend that it cover accounts error-theoretic, pragmatic, and whatever else in between.

⁷ Here’s a general complaint. Certain of the contrasts in the data we have seen (and will see) look especially difficult to explain along pragmatic or error-theoretic lines. For example: as per §2.3’s discussion of knowledge-wh versus knowledge-that, our intuitive judgments about the *History Exam* ascriptions seems to depend sensitively on whether the complement to ‘know’ is a wh-clause or a that-clause. This sensitivity looks to be difficult to square with the loose speech theorist’s account of our intuitions. According to this theorist, the reason we find assertions of (1)–(6) felicitous is because we associate certain kinds of non-literal contents with those speech acts, and judge that *those* contents are true. But if a fine sounding assertion of (6, ‘Peggy knows *what* the answer to the first question is’) is really just loose speech used to convey a true proposition other than the one that is (6)’s semantic content, then why does an assertion of (6*, ‘Peggy knows *that* the answer to the first question is b’) tend to sound so much worse? Is it really so much harder to convey the intended message in the language of knowledge-that? It is difficult to imagine what would explain this. More plausible is the idea that something lexical is responsible for the contrast, rather than something like a glitch in our capacities for mind-reading.

⁸ Though see Hazlett (2010, 2012) for arguments against orthodoxy.

For instance: it is well established that sometimes when we wish to describe or predict the behavior of other agents, we use sentences that portray the way the world *appears* to those agents, regardless of whether those appearances are veridical.⁹ To use just one example: if you ask me why my Flat-Earther friend won't let his child visit space.com, I can felicitously answer 'Because it's NASA propaganda', even when it is clear to both us that no such thing is true. It's just a convenient way of expressing how the world looks from my friend's perspective. Knowing that there is a general "perspective-shifting" phenomenon, we should expect there to be contexts in which the facts about what agents know and the facts about what those agents *take* themselves to know are treated interchangeably. Consequently, we can invoke a general mechanism to explain way the existence of felicitous non-factive uses of 'know'. And in doing so we can leave the semantic validity of the inference from 'S knows that p' to p unscathed.

I know of no comparably general account of what is driving our intuitive judgments about cases like *History Exam*, *Office Meeting*, and *Memory Experiment*.¹⁰ Until one is offered, I see little reason to think that treating the judgments non-semantically has any explanatory advantage over treating them semantically. I will thus continue to assume that our judgments about these cases ought to be taken at face value, rather than be quarantined as a pragmatic oddity or error in conceptual competence. And with that, I now turn to developing a theory of knowledge that can make sense of them.

3 Knowledge by constraint

What must a theory of knowledge look like if it is to account for the strange ascriptions of *History Exam*?

3.1 Contextualism versus invariantism

For starters, it will have to be some version of *contextualism* about knowledge—the view that 'know' is associated with different epistemic relations across different occasions of use. The alternative is to adopt a kind of *ultra-liberal invariantism* about knowledge, claiming that knowledge is in every context as liberal as Peggy's knowledge of the answer to the first test question.¹¹

Aside from its intuitive implausibility, such an extreme form of invariantism faces a number of empirical problems. For starters, it is relatively easy to get *bad* readings of just about all of (1)–(6). We saw in §2.3 that changing the complement of 'know' from a *wh*-clause to the a *that*-clause tends to have this effect. But there are simpler ways to achieve it too. The same

⁹ See, e.g., [Schlenker \(2004\)](#) for background on this phenomenon.

¹⁰ Note that perspective-shifting is *not* a plausible diagnosis of our intuitions about these cases. Peggy's perspective is the same as ours in all the respects that matter.

¹¹ Here I set aside relativism about knowledge and ignore the difference between moderate- and sensitive invariantism. Complications that might arise from considering these sorts of options will not affect the main points of this section. But see the discussion of "psychologism" in the following section for some reasons to be skeptical of sensitive invariantist accounts of the data.

can happen when we start asking questions like ‘Does Peggy *really* know what the answer to the first question is?’. Or when we make a big deal of the fact that, though Peggy may *in fact* have the good answers, she doesn’t *know* that she does. So even though it is easy to hear (1)–(6) as having good readings, it is also easy to hear them as having bad readings. This level of malleability in our intuitive judgments is difficult to explain on the assumption that these ascriptions are context-invariantly true.

There are also some more exotic constructions that pose direct challenges to invariantist treatments of the data. For example, in the context of *History Exam*, Peggy could naturally think or say something along the lines of:

- (12) If I have the good answers, then I know that unbeknownst to me or Pete, Pete is going to fail the exam.

But assuming with the invariantist that ‘know’ only ever expresses one epistemic relation, (12)’s consequent entails a sentence of the form ‘I know that: p and I don’t know that p’, which in turn entails the straightforwardly contradictory ‘I know that p and I don’t know that p’. So this cannot plausibly be the natural reading of the sentence. Instead, the natural reading of (12) must involve a non-uniform interpretation of the two occurrences of ‘know’: the outer ‘know’ expressing the kind of strange knowledge brought out by the earlier examples, the inner ‘know’ expressing a more ordinary kind.

Thus, the real question is what form of contextualism about knowledge is best suited to explain what is happening with the data under consideration. Might any existing contextualist theory of knowledge be up to the task?¹² The answer to this question depends on how strictly these views are held to their original letter. If held strictly enough, I’m skeptical that any will get things right. But I lack the space to give a proper argument for this; or to explore, for each view, whether and how it needs to be revised to give so as to model the data correctly.

So instead the plan is this. This section will outline a bare-boned, contextualist semantics for ‘know’ in the style of Lewis’s (1996) contextualist relevant alternatives theory. I’ll start with a carbon-copy of his view, but then shortly after identify an important technical adjustment needed to give an adequate treatment of *embedded* knowledge ascriptions like (1)–(4). Crucially, the presentation will be entirely silent on the *meta-semantics* of the proposed semantic theory—that is: it will be silent on issues about how, when, and why the denotation of ‘know’ shifts across contexts. The section that follows this one will take up that problem.

¹² For a sampling of such theories, see e.g., Cohen (1986, 1988, 1999); DeRose (1992, 1995); Lewis (1996); Heller (1999); Neta (2003); Schaffer (2005b, 2007); Blome-Tillmann (2009); Ichikawa (2011); Schaffer & Szabo (2014).

3.2 The semantics

On a *relevant alternatives* conception of knowledge, to know that p is to eliminate all the relevant alternatives to p .¹³ Here is a contextualist-friendly regimentation of this slogan:

R.A. SCHEMA ‘S knows that p ’ is true at world w and context c ($\langle w, c \rangle$) iff: at w , all the c -relevant worlds are p -worlds.¹⁴

Allowing myself some looseness with use and mention, the intuitive idea behind R.A. SCHEMA is this. In any given context there is a set of worlds against which a claim to knowledge is tested—namely the worlds that are relevant in context. For a world to be relevant in context, it has to be compatible with two bodies of information: first, a body of information associated with the *subject* of the ascription; and second, a body of information associated with the *context* of the ascription. Only if a world is compatible with both bodies of information is it relevant. And if every world that is relevant is a p -world, then by lights of that context the subject knows that p .

What is the body of information associated with the subject of a knowledge ascription? It’s the subject’s **evidence**, which I’ll model as a set of worlds (i.e., a proposition). What is the body of information associated with the context of the ascription? A **restriction**, which I’ll also model as a set of worlds (i.e., a proposition). What a subject knows in context is whatever is entailed by the intersection of the subject’s **evidence** and the contextually supplied **restriction**. Being less loose with use and mention: ‘S knows that $-p$ is true at $\langle w, c \rangle$ just in case: all the worlds compatible with both S’s **evidence** (at w) and the **restriction** (as determined by c) are p -worlds.

Now to say a bit more about how **evidence** and the **restriction** are working. Officially I will leave **evidence** as an unanalyzed primitive, assuming only that it represents a reflexive epistemic relation, and thus that an agent’s **evidence** at a world w always contains w . This is to secure the factivity of knowledge. For the purposes of explaining the view, however, it will be helpful to give a working theory of **evidence** in terms of a more familiar epistemic notion. So in an exploratory spirit I will identify one’s **evidence** with one’s *ordinary knowledge*.¹⁵ I will use “*knows*” (that is, the *asterisks*) to denote ordinary knowledge and related notions (e.g., *epistemic possibility*, *ignorance*, etc.). The basic idea is that *knowledge* is what the typical moderate invariantist about knowledge thinks knowledge is. I will thus assume that facts about *knowledge* are context-invariant. I will also assume that most people *know* whether they have hands, who the current president of the United States is, where they’ll be living in a week, and so on. I

¹³ Proponents of the relevant alternatives conception of knowledge include Dretske (1970); Unger (1975); Lewis (1996); Schaffer (2005a,b); Rysiew (2006); Blome-Tillmann (2009); Ichikawa (2011); Schaffer & Szabo (2014).

¹⁴ For the sake of concreteness, I will assume that the objects of knowledge are propositions, and that propositions are sets of possible worlds. This is simply a modeling choice. The important features of the account (as designed to handle cases like *History Exam*) are compatible with a wide variety of views on the nature of content.

¹⁵ This is a departure from Lewis, who uses something along the lines of *Cartesian certainty* in spelling out the **evidence** role. Since we are not (yet) interested in accounting for skeptical puzzles, *ordinary knowledge* is more convenient for present purposes. I will return to these issues in §6.2.

will further stipulate that *knowledge* satisfies CREDENCE, BELIEF, and JUSTIFICATION. Thus, *History Exam*'s Peggy does *not* *know* what the answer to the first test question is.

That covers **evidence**: it's a person's (ordinary) *knowledge*. The idea behind the **restriction**—the set of worlds supplied by context—is again easiest to grasp given some looseness with use and mention. Intuitively, we want the **restriction** to represent the conjunction of all the propositions the subject gets to “know for free” in context. That is: the strongest proposition one gets to know without having any **evidence** for. In some contexts the restriction may be something trivial like the tautology \top —in which case agents will only know what is entailed by their **evidence** (i.e., what they *know*). In others it may be something weightier, say the proposition Ψ , in which case agents will know any p such that their **evidence** entails the material conditional ($\Psi \supset p$). But the basic idea is that what changes from context to context is how much work an agent's **evidence** (*knowledge*) is put to in establishing what they know.

However, for reasons that will emerge shortly, rather than have the value of the **restriction** be supplied *directly* by context, I will have context supply a *question*—i.e., a partition over possible worlds¹⁶—from which the **restriction** is derived. More precisely, I will follow [Dorr & Hawthorne \(2013\)](#) in using the mechanisms of *constraint* to generate the contextually variable **restriction**.¹⁷ Constraints are polar questions (i.e., yes/no questions). Constraints combine with worlds to generate propositions, which can then be used to restrict sets of worlds (like a subject's **evidence**) via set intersection. The proposition a constraint question generates at a world is just the true answer to that question at that world. So, supposing $\Psi^?$ is the constraint \ulcorner Is it true that $\Psi?$ \urcorner , and that Ψ is true at w but false at w' , then the proposition generated by $\Psi^?$ at w is Ψ , while at w' is $\neg\Psi$. Thus: context's contribution to the calculation of the interpretation of 'know' will go by way of a constraint question, not a proposition; and the **restriction** will be whichever proposition the relevant world says is the true answer to the contextually relevant constraint question.

Combining a *knowledge*-based analysis of **evidence** with a constraint-based analysis of **restriction** yields the following implementation of R.A. SCHEMA:

BASIC CONSTRAINT 'S knows that p ' is true at $\langle w, c \rangle$ iff: every world compatible both with S's *knowledge* (at w) and with the true answer (at w) to the c -determined constraint $\Psi^?$ is a p -world.¹⁸

¹⁶ See, e.g., [Hamblin \(1958\)](#), [Groenendijk & Stokhof \(1984\)](#).

¹⁷ [Dorr & Hawthorne \(2013\)](#) use the mechanism to explain various puzzling data involving epistemic modals. A full discussion of the relation between the embedded uses of 'know' of interest to this paper and the embedded uses of modals of interest to theirs is beyond the scope of the present discussion. But readers familiar with their work should find the similarities striking. See also [Hawthorne \(2007\)](#) for more on the kind of modal data that is thought to exhibit the constraint phenomenon, as well as [Gillies \(2010\)](#) and [Moss \(2015\)](#) for other examples of *constraint-like* accounts of the semantics of epistemic vocabulary.

¹⁸ Apropos of the previous footnote, **BASIC CONSTRAINT** finds semantic analogies between 'know' and other expressions that are subject to the mechanisms of constraint. If [Dorr & Hawthorne \(2013\)](#)'s claims are right (in broad outline), then epistemic modals like 'must' and 'certainly' are the closest semantic analogs to 'know'. This makes 'know' a kind of adverbial quantifier (like 'always' or 'must') rather than a domain quantifier (like 'all' and 'every'—cf. [Ichikawa \(2011\)](#)). See [Schaffer & Szabo \(2014\)](#), §§2.4, 3.1, and 4.5 in particular for discussion of the advantages modeling 'know' along the lines of a quantificational adverb (as **BASIC CONSTRAINT** does), particularly as concerns standard

In words: if we want to know whether ‘S knows that p’ is true at $\langle w, c \rangle$, we look at all the worlds compatible with what S *knows* at w , prune off those that are inconsistent with the true answer to the c -determined constraint $\Psi^?$, and see if any of the worlds that remain are ones in which $\neg p$. ‘S knows that p’ expresses a true proposition if and only if none of those worlds are $\neg p$ -worlds. A more illustrative but less precise way of thinking about it: if ‘S knows that p’ is true when the c -determined constraint is $\Psi^?$, then if S were to come to *know* the true answer to $\Psi^?$, S would thereby be in a position to *know* p too.

3.3 Accounting for the data

Now to apply BASIC CONSTRAINT to *History Exam*. Recall that I am making no assumptions about the meta-semantics of constraint, and so will help myself to whatever resolution of the context is needed to get the data right. The aim here is to illustrate certain structural features a theory of knowledge needs to have if it wants to account for the relevant readings. I will return to the question of whether there might be principled ways of deriving these readings in detail in the next section of the paper.

Here again are *History Exam*’s (1)–(6):

[Peggy (student):]

- (1) If my answer sheet is good, then I know what the answer to the first question is.
- (2) One of us knows what the answer to the first question is.
- (3) Either Pete got the good answers, or I know what the answer to the first question is.
- (4) I might know what the answer to the first question is; it depends on whether I got the good answers or not.

[Roger (TA):]

- (5) Whoever has the answer sheet that says ‘1: b’ knows the answer to that question.
- (6) Peggy knows what the answer to the first question is.

According to BASIC CONSTRAINT, ‘Peggy knows that the answer to the first question is b’ is true at $\langle w, c \rangle$ just in case every world compatible with (i) what Peggy *knows* at w and (ii) the true answer to the c -determined constraint is a world in which the answer to the first question is b. Let G be the proposition that Peggy’s answers are good, and let $G^?$ be the corresponding constraint question (to which, at w , G happens to be the true answer). Assuming that the natural readings of (1)–(6) are those on which the c -determined constraint is $G^?$, we get that: ‘Peggy knows that the answer to the first question is b’ is true at $\langle w, c \rangle$ just in case every G -world compatible with what Peggy *knows* is a world in which the answer to the first question is b. This is indeed the case. Since Peggy *knows* that *if* her answers are good then the answer to

semantic challenges to epistemic contextualism.

the first question is b, it follows that there are no G-worlds compatible with what Peggy *knows* in which the answer to the first question isn't b. Thus 'Peggy knows that the answer to the first question is b' is true at $\langle w, c \rangle$.¹⁹

It is straightforward to see how the *truth* of 'Peggy knows what the answer to the first question is' entails the truth of all of (1)–(6) (given the details of *History Exam*). But there remains the question of why, intuitively, Peggy's has to *hedge* when she ascribes herself knowledge of the answers by embedding that ascription under a conditional, disjunction, modal, etc., whereas Roger can felicitously ascribe Peggy that same knowledge unembedded with an ascription like (6). As it turns out, the answer to this question will also explain why BASIC CONSTRAINT generates its **restricting** propositions through constraint questions.

But first some new notation. I will say that 'know_Ψ' denotes whatever relation BASIC CONSTRAINT predicts 'know' denotes when the constraint question is Ψ'. Thus: S knows_Ψ that p iff the intersection of S's *knowledge* and the true answer to Ψ' entails p.

Now for the explanation. BASIC CONSTRAINT's account of why Roger can assert the unembedded (6) while Peggy has to hedge by embedding it under a conditional, quantifier, disjunction, or modal is simple. It is that only Roger *knows* that Peggy knows_G what the answer to the first question is. Why does only Roger *know* that Peggy knows_G what the answer to the first question is? Because only Roger *knows* that the answer to G' is G. For all Peggy *knows*, the answer to G' is ¬G. Thus, for all Peggy *knows*, she knows_G that the answer to the first question is *not* b. This is why only Roger can felicitously assert (6).

What Peggy does *know*, however, is that she knows_G what the answer to the first question is *conditional on G being true*. Thus, she *knows* that *if* G, then she knows_G what the answer to the first question is. Hence why she can assert (1). Likewise, she *knows* that it *might* be that G, and thus *knows* that it might be that she knows_G what the answer to the first question is. Hence why she can assert (4). And the same points apply *mutatis mutandis* to (2) and (3). Thus, constraint questions play an essential role in accounting the asymmetry in the intuitive judgments about what Peggy and Roger may felicitously *assert* in the context of *History Exam*.

As it turns out, they also play an essential role in accounting for the truth-conditions of certain other embedded ascriptions. For here is another thing Peggy think or say in the context of *History Exam*:

(13) There's a 50% chance I know what the answer to the first question is.

(13) has a true reading. It can't be one on which 'know' denotes *knowledge*, for Peggy *knows* that she doesn't *know* what the answer to the first question is. So it must take a restricted reading of some sort. However, if the **restriction** were generated directly by a proposition—say, G—then (13) would still be false. And that's because there is a 100% chance (not 50%) that

¹⁹ Note that the other student, Pete, does *not* *know* that if Peggy's answers are good, then the answer to the first question is b. This means that 'Pete knows what the answer to the first question is' is false in any context in which the constraint is G'—a desirable prediction.

all the worlds consistent both with what Peggy *knows* and with the proposition that Peggy has the good answers are worlds in which the answer to the first question is b. To get it so that there is only a 50% chance that the proposition expressed by ‘I know what the answer to the first question is’ is true, it needs to be that Peggy is only 50% certain about whether she stands in the relevant epistemic relation to the proposition that the answer to the first question is b. And that is exactly what is achieved by using a constraint question like $G^?$ to generate the **restriction**, rather than reading it off from context directly. Peggy is 50% certain that the answer to $G^?$ is G. She is thus 50% certain that every world compatible with what she *knows* and with the true answer to $G^?$ is a world in which the answer to the first question is b. This is why she may truly assert (13).

I contend that BASIC CONSTRAINT provides a simple and intuitive semantic account of the behavior of ‘know’ in *History Exam*’s (1)–(6). Working in the tradition of Lewis (1996), the core idea is that ‘know’ expresses a relation between an agent and all the propositions that are entailed by a contextually determined subset of that agent’s *knowledge*. The relevant subset is determined by the mechanisms of constraint—the flexibility of which explains why *History Exam*’s Peggy can be ascribed knowledge of propositions for which she possess at best dubious evidence, and the structure of which explains why Roger may ascribe Peggy this knowledge unconditionally even when Peggy herself cannot.

4 The meta-semantics of constraint

In explaining BASIC CONSTRAINT’s account of *History Exam*, I assumed without argument that, on the relevant (true) readings of (1)–(6), the contextually supplied constraint question is something along the lines of ‘Does Peggy have the good answers?’. One might wonder whether this feature of context can be predicted (or at least explained), rather than simply reverse-engineered to fit the intuitive judgments.

This is a question about the *meta-semantics* of ‘know’: namely, the question of what explains when and why context favor some interpretations (i.e. constraints) rather than others. With respect to this question, I will begin by arguing that the data considered so far raise significant problems for standard contextualist meta-semantics for ‘know’ (§4.1). I will then argue that given various idiosyncrasies in the intuitive judgments, it is hard to see how there could be *any* predictive meta-semantics for ‘know’ (§4.2). But against taking this to be reason to dismiss a contextualist treatment of the data along the lines of BASIC CONSTRAINT, I will also argue that the difficulties involved here are problems for everyone, and in fact may be a symptom of the conceptual primitiveness of knowledge (§4.3).

4.1 Against standard contextualist meta-semantics

There is a certain meta-semantic assumption common to most existing forms of contextualism about knowledge. Put roughly, it is the assumption that contextually driven shifts in the interpretation of ‘know’ are due to shifts in the high-level psychological states of those making the relevant knowledge ascriptions. See, e.g., [Cohen \(1999\)](#):

...[S]entences of the form ‘S knows p’ can, at one time, have different truth-values in different contexts. Now when I say ‘contexts’, I mean ‘contexts of ascription’. So the truth-value of a sentence containing the knowledge predicate can vary depending on things like *the purposes, intentions, expectations, presuppositions, etc.*, of the speakers who utter these sentences. (p. 1; my emphasis)

As well as [Rysiew’s \(2016\)](#) survey article:

...[T]here is a strong degree of consensus among the theories under discussion²⁰ that context itself is to be understood in terms of such things as *the interests, purposes, expectations, and so forth*, of the knowledge attributor. (§3.3; my emphasis).

Let X denote one’s preferred admixture of the factors Cohen and Rysiew allude to.^{21,22} We can then state the core meta-semantic assumption common to most existing forms of contextualism about knowledge as follows:

PSYCHOLOGISM Changes in the semantic contribution of ‘know’ across contexts are to be explained in terms of changes in the X facts.

Though PSYCHOLOGISM is vague, it can be shown that any reasonable precisification of it does poorly with the various embedded uses of ‘know’ raised to salience by *History Exam*.

For starters, our judgments about what Peggy can say or think in *History Exam* seem to be exactly the opposite of what we would expect taking the X facts at face value. Take (4) (‘I might know what the answer to the first question is; it depends on whether I got the good answers or not’), for instance, and consider, in turn, Peggy’s expectations, presuppositions, and interests. Peggy is clearly taking seriously the possibility that she has bad answers—in fact, she is literally announcing its possibility. So it’s not like she is expecting or presupposing that she has the good answers. And it is a matter of enormous practical significance to Peggy that she answers the

²⁰ Here the theories referenced include those of [Heller \(1995\)](#); [Lewis \(1996\)](#); [Cohen \(1998\)](#); [Rieber \(1998\)](#); [Neta \(2002\)](#); [Blome-Tillmann \(2009\)](#); [Greco \(2010\)](#).

²¹ It won’t matter for this section’s arguments whether the X factors are ascriber-, subject-, or assessor-dependent.

²² One potential source of “ X facts” not mentioned in either the Cohen or Rysiew quotations: facts about the presuppositions of the “question under discussion” (QUD), in the style of (e.g.) [Ginzburg \(1996\)](#) and [Roberts \(2012\)](#). QUDs are what [Schaffer & Szabo \(2014\)](#) use to make substantive predictions about how the context-sensitivity of ‘know’ gets resolved. The arguments of this section apply just as much to the QUD-theoretic meta-semantics as to the interest/attention/presupposition-theoretic ones. See also [van Elswyk \(2020\)](#) for skepticism about QUD-based meta-semantics for ‘know’.

test questions correctly—in fact, we could make it a life or death matter without disrupting the relevant intuitions. So ‘know’ should simply not be amenable to a reading on which ‘Peggy knows what the answer to the first question is’ expresses a truth in Peggy’s mouth, whether embedded under a modal or not. Yet our judgments about (4) suggest otherwise.

Another problem for PSYCHOLOGISM is that the natural readings of many of *History Exam*’s ascriptions involve *intra-sentential* shifts in the denotation of ‘know’, shifts that cannot plausibly correspond to anything like a change in the X facts. Consider (12, ‘If I have the good answers, then I know that unbeknownst to me or Pete, Pete is going to fail the exam’). On pains of (attitudinal) inconsistency, ‘know’ must express two different relations within the scope of (12)’s consequent. But if PSYCHOLOGISM is true, then changes in the denotation of ‘know’ can occur only when, for the relevant parties, the X facts change. To get the natural reading of (12), then, we would have to hold that changes in what’s relevant, salient, at stake, presupposed, etc., occur *mid-consequent*. This is not a desirable result. To name just one implausible consequence: it would seem to suggest that one of the words in the string ‘that unbeknownst to me or Pete, Pete is going to fail the exam’ could mark the occasion whereby the X facts shift from those amenable to a liberal reading of ‘know’ to a more ordinary one.

A similar point can be made about simpler embedded cases like (1, ‘If I have the good answers, then I know what the answer to the first question is’). Think of things from Peggy’s perspective. Either the high-level features of your psychology responsible for settling the interpretation of ‘know’ (e.g., what you’re attending to, or what your presuppositions, interests, etc., are) are such that ‘I know what the answer to the first question is’ expresses a true proposition in your mouth, or they are such that it expresses a false proposition. If they are such that it expresses a *true* proposition, then what explains your first-order uncertainty about what the answer to the first question is? *Some* feature of your psychology is causing you to (subjectively) assess it to be no more than .5 likely that the answer to that question is b. Given PSYCHOLOGISM, shouldn’t this be exactly the kind of thing that induces a reading of ‘know’ on which ‘I know what the answer to the first question is’ expresses a falsehood?

Let us assume so: that is to say, let us conclude from the foregoing that the sentence ‘I know what the answer to the first question is’ expresses a *false* proposition in your mouth prior to the utterance of (1). Now a new problem arises. (1) is true. So is its antecedent. So the proposition expressed by ‘I know what the answer to the first question is’ in consequent position must be other than the one it expresses unembedded. Since ‘I know what the answer to the first question is’ expresses a false proposition in your mouth prior to the utterance of the conditional, it follows that the context shifts *mid-conditional* so as to make it express a true proposition in consequent position. But, as in the case of (12), one wonders what kind of psychological mechanisms could realistically underly such a shift.

As I see it, we do better not making the assumption that generates these problems. They arise because standard contextualist theories of ‘know’ (and their corresponding meta-semantics)

were designed to account for the pull of skepticism in the face of lottery puzzles, bank cases, and skeptical paradoxes. These are cases in which the X facts can quite plausibly be taken to explain the shifts in the intuitive judgments about the extension of ‘know’. But these theories were not designed for constructions like (1)–(4) or (12). What these constructions show is that ‘know’ can undergo fine-grained, intra-sentential shifts in denotation—shifts that are not exclusively responsive to changes in the high-level psychological states of those making or assessing claims that contain it. Thus, setting aside the question of whether BASIC CONSTRAINT is true, the fact that there are true readings of sentences like (1) and (12) shows that PSYCHOLOGISM must be false.

This is an interesting result in its own right. But it also makes the task of pairing BASIC CONSTRAINT with a predictive meta-semantics that much more difficult. If not facts about attention, stakes, presuppositions, etc., then what *does* explain why the constraint takes the value it does on the natural readings of the ascriptions in cases like *History Exam*?

4.2 Making the data even stranger

I am not sure. In fact, I think it is unclear whether we should expect any satisfying answer to this question. This is not just because the judgments about cases like *History Exam* are hostile to PSYCHOLOGISM; more fundamentally, it is because certain patterns in the data seem resistant to being modeled by *any* kind of predictive theory.

One case in point concerns the contrasts (discussed in §2.3) between knowledge-that and knowledge-wh ascriptions. Given the details of *History Exam*, we can be confident that any context in which ‘Peggy knows what the answer to the first question is’ has a true reading is one in which ‘Peggy knows that the answer to the first question is b’ has a true reading, and vice-versa. This is why (8) seems like a contradiction:

- (8) ?? Although Peggy knows what the answer to the first question is, she does not know that the answer to the first question is b.

Nonetheless, the contexts most naturally evoked by knowledge-that ascriptions are ones on which (1)–(6) express falsehoods, while the ones most naturally associated with the knowledge-wh are ones on which (1)–(6) express truths. But why should there be this meta-semantic difference given the seemingly obvious (contextual) entailments between the two kinds of knowledge ascriptions? What general psychological mechanism would explain such a thing?

Another case in point. A natural meta-semantic hypothesis to pair with BASIC CONSTRAINT is that for conditional knowledge ascriptions like (1), the contextually supplied constraint is just the antecedent asked as a polar question. This would at least explain all the conditional ascriptions considered so far. But it doesn’t take much work to show that that can’t be right—at least not in full generality—for consider (1)_T’s “tautologous” analog:

- (1)_T ? If my answer sheet is good, then I know whether my answer sheet is good.

It is straightforward to see why (1_T) is true in any context in which the constraint question is ‘Does Peggy have the good answers?’—exactly the kind of context in which I’ve argued (1) expresses a truth. But to my ears (1_T) is noticeably degraded. This seems to be part of a general pattern. Consider, e.g.:

(14) ✓ If Biden has had a fatal heart attack in the past few moments, then I know Harris is now the president of the United States.

(14_T) ? If Biden has had a fatal heart attack in the past few moments, then I know Biden has had a fatal heart attack in the past few moments.

(14) seems fine; (14_T) does not. Yet presumably any context in which the former is true the latter is as well. So perhaps the meta-semantic generalization ought instead be something along the lines of: for conditional ascriptions, the contextually supplied constraint is the antecedent asked as a polar question—*so long as the constraint doesn’t “trivialize” the embedded knowledge ascription.*

Indeed, there is evidence that natural language imposes this kind of meta-semantic constraint on the interpretation of adverbial quantifiers more generally.²³

(15) ✓ If I play poker, I always wear my lucky hat.

(15_T) ? If I play poker, I always play poker.

(16) ✓ If Paul is in Paris, it is certain he is in France.

(16_T) ? If Paul is in Paris, it is certain he is in Paris.

So that is something. But how exactly is one supposed to understand the relevant notion of ‘trivialize’? Taking (16) as a cue, consider:

(17) a. ✓ If Paul is in Paris, then I know he’s in France.

b. ✓ If Paul is in Paris and Paris is not Berlin, then I know he’s not in Berlin.

These both seem fine, even despite the fact that it is about as trivial as it gets that (e.g.) if Paul is in Paris and Paris is not Berlin, he’s not in Berlin. So evidently there is something *particularly* trivial about conditional ascriptions of the form ‘If p, then I know that p’. But it is not at all clear how to predict this fact given standard theoretical tools.

And still there is the most basic worry for the prospects of a predictive meta-semantics: how are we explain the fact that in the vast majority of contexts of interest to epistemologists, principles like CREDENCE, BELIEF, and JUSTIFICATION seem to express truths, yet when we embed knowledge ascriptions under conditionals, quantifiers, disjunctions, and modals (as in *History Exam’s* (1)–(4)), counterexamples seem to arise with ease?

²³ Cf. Schaffer & Szabo (2014, §4.3).

To add another example to the pile, many theorists have taken [Gettier \(1963\)](#) to have shown decisively that knowledge is incompatible with the truth of one's belief being a matter of a certain kind of luck. But cases like *History Exam* provide *prima facie* evidence that in at least some contexts, Gettier's generalization fails. Moreover, we can make the same point using one of the standard Gettier cases:

Standard Gettier Smith's boss Roger tells Smith that Jones is getting a promotion. Smith knows Jones has ten coins in his pocket, so trusting Roger's testimony he infers that the man who will get the promotion has ten coins in his pocket. But unbeknownst to Smith, Roger is confused: it is Smith, not Jones, who is getting the promotion. Also unknown to Smith is that he himself has ten coins in his pocket. So Smith's belief that the man who will get the promotion has ten coins in his pocket is true after all.

It is widely thought that in the contexts most naturally evoked by *Standard Gettier*, (18) expresses a falsehood:

(18) Smith knows that the man who will get the promotion has ten coins in his pocket.

But putting ourselves in Smith's shoes, we can just as well say and think things like:

(19) If Roger confused me with Jones and I happen to have ten coins in my pocket, then I know whether the man who will get the promotion has ten coins in his pocket.

(20) I might know how many coins are in the pocket of the man who will get the promotion; depends on whether I have ten coins in my pocket and Roger confused me with Jones.

Given that Roger confused Smith with Jones and that Smith happens to have ten coins in his pocket, it looks to follow that (18) has true readings after all. So how are we to reconcile the fundamental Gettier intuition with the ease with which constructions like (19)–(20) can be rattled off? Again, it is not clear.

4.3 What now?

In light of these considerations, ought one conclude that there can be no non-trivial meta-semantics for 'know'? Probably not. These constructions have received little attention in the literature, so it is anyone's guess what future theorizing will bring to bear. But as it stands I do not have a theory that says in general terms when 'know' will take one interpretation rather than another. And one might take this to be strong evidence that there is something wrong with a contextualist semantics for 'know' along the lines of BASIC CONSTRAINT.

I want to spend the rest of the section pushing back against this popular reaction. I think there is no good reason to think that a contextualist semantics for an expression ought to be accepted only if we've found a predictive meta-semantics for it. That is to say: I think there

is no good reason to think that we need a theory of the mechanisms that determine when, how, and why context shifts with regards to some expression to know (i) that that expression is context-sensitive and (ii) that it undergoes certain kinds of shifts. Take quantifier domain restriction, for instance. We can know that the domains over which expressions like ‘everyone’ and ‘the sailors’ quantify vary from context to context—sometimes the universe, sometimes just the people on this boat—even if we don’t have a general theory that tells us when these shifts happen and why. To reject a contextualist semantics for quantifiers because it has not been paired with a sufficiently predictive meta-semantics for domain restriction is to confuse importantly distinct theoretical questions. One should not make a similar mistake in assessing the prospects of contextualist accounts of ‘know’.²⁴

For the sake of argument, however, let us suppose that we will simply never find a predictive meta-semantics with which to pair BASIC CONSTRAINT. Given the flexibility of the theory, none can be given. *Now* do we have reason to reject it? It is hard to see why we would, unless we took the difficulty of finding a predictive meta-semantics to be evidence that the data ought not be modeled semantically at all. But what would justify *that* inference? The patterns in our intuitive judgments about knowledge ascriptions do not suddenly become less mysterious just because we decide that some of them must be in error or the byproduct of a yet to be specified pragmatic mechanism. For there remain the questions of what explains when and why these errors arise; or when the pragmatic mechanism gets put in use, what its rules are, and how it interacts with general features of communication and rationality. And it’s not like we have good reason to believe we will do any better with *those* questions than with the ones about the contextualist’s meta-semantics. The fact that the knowledge ascriptions surveyed in this paper seem resistant to rational explanation is a problem for everyone, not just those who take the judgments at face value. Until we’ve convinced ourselves that we can give a better explanation of what’s happening by not treating them as semantically respectable—as in (e.g.) the “perspective shifting” treatment of non-factive uses of ‘know’—we should be prepared to follow the intuitions where they lead. If that means readjusting our expectations with respect to our ability to predict when a knowledge ascription will seem true, then so be it.

In relation to this point, consider the thesis popularized by [Williamson \(2000\)](#) that knowledge is conceptually primitive. Let us suppose Williamson is correct: though there may be various interesting generalizations we can make about the proper use of the concept of knowledge, we cannot know in general and non-circular terms the conditions for its application. Given this assumption, we should also think we cannot know in general and non-circular terms the conditions under which an utterance of ‘S knows that p’ will express a true proposition. Else we would be in the curious position of knowing in general and non-circular terms the conditions under which an utterance of ‘S knows that p’ expresses a true proposition, but not knowing in general and non-circular terms the necessary and sufficient conditions for knowing. However,

²⁴ [Ichikawa \(2011\)](#) makes a similar point in defending [Lewis’s \(1996\)](#) contextualist semantics from various complaints about its predictiveness.

if we cannot know in general and non-circular terms the conditions under which an utterance of ‘S knows that p’ expresses a true proposition, then it is not clear why we should expect to be able to give a non-trivial meta-semantics for ‘know’. After all, the opacity of its meta-semantics may well be what *explains* why we cannot know in general and non-circular terms the conditions under which ‘S knows that p’ expresses a true proposition. Those convinced of the conceptual primitiveness of knowledge are already committed to the view that our intuitions about knowledge form something of a black box. Unless they have some special argument that it has to be the semantics *and only the semantics* that grounds this fact, they should not be resistant to the possibility that the meta-semantics is at least partially responsible too.

5 From BASIC CONSTRAINT to THINKING CONSTRAINT

In the absence of a meta-semantic theory that gives the conditions under which context supplies one constraint question rather than another, the proponent of BASIC CONSTRAINT places very few necessary conditions on the existence of a true reading of ‘S knows that p’. In fact, the only necessary condition is that p be true. For if p is true, then so long as there is a true proposition Ψ such that S *knows* that $\Psi \supset p$, then ‘S knows that P’ will be true whenever the constraint is Ψ ?

Despite having just argued that we should be prepared to have very little to say about when ‘know’ takes one reading rather than another, I believe there is *an* amendment that ought to be made to BASIC CONSTRAINT so as to make it more theoretically satisfying. This is because there is a certain class of knowledge ascriptions that seem to have no true readings at all, yet ought to have true readings given the combination of BASIC CONSTRAINT and the non-existent meta-semantics. And there is a simple generalization that can be incorporated into the view so as to avoid this problem—one that raises some interesting questions about the relationship between the concepts of knowing, thinking, and believing.

5.1 Two contrasts

Recall the basic details of *History Exam*. Roger sells answer sheets to Peggy and Pete. One of them gets bad answers, but neither can figure out who it is in time for the test. All they know is that their answer sheets differ on every question. Eventually Roger discovers it was Pete who got the bad answers, and upon doing so thinks to himself:

- (6) ✓ Peggy knows what the answer to the first question is.

We saw in §3.3 why given the constraint question G ? (‘Does Peggy have the good answers?’) BASIC CONSTRAINT can account for the true reading of (6): (Roger *knows* that) Peggy *knows* that if she has the good answers, then the answer to the first question is b; Peggy has the good answers; so all the worlds compatible with what Peggy *knows* and with the true answer to G ?

are worlds in which the answer to the first question is b; so she knows_G what the answer to the first question is.

What is of interest is that Roger cannot seem to felicitously think or say (21) (at least not unembedded):

(21) ? Peggy knows who has the good answers.

We have been given no reason to expect this. According to BASIC CONSTRAINT, whenever the constraint question is something like ‘Is the answer to the first question b?’ (hereafter 1[?]), (21) should have a true reading. And this is because Peggy *knows* that if the answer to the first question is b, then she has the good answers. Indeed, Roger *knows* that Peggy *knows* this. Since the answer to the first question is in fact b, it follows that all the worlds compatible with what Peggy *knows* and the true answer to 1[?] are worlds in which Peggy has the good answers, and thus that Peggy knows₁ who has the good answers. So either there is some mysterious feature of the meta-semantics of ‘know’ that explains why G[?] but not 1[?] is a natural constraint question in the contexts associated with *History Exam*, or the contrast between (6) and (21) is evidence that BASIC CONSTRAINT is in need of amendment.

That’s the first unexplained contrast. To see the second, I’ll add the following detail to *History Exam*: suppose the reason Roger doesn’t know who got the good answers is not because he forgot to whom he gave which set, it’s because he himself doesn’t know which set of answers is the good one. That is to say: Roger knows he gave Peggy answer sheet 1 and Pete answer sheet 2, and he knows what each answer sheet says, but he has no idea which has the correct answers. (And suppose Peggy knows all this.)

On this version of the case, Peggy can continue to felicitously think or say things like:

- (1) ✓ If I have the good answers, then I know what the answer to the first question is.
- (4) ✓ I might know what the answer to the first question is; it depends on whether I got the good answers or not.

But observe that Peggy cannot seem to felicitously think or say things like:

- (22) ? If I have the good answers, then Roger knows what the answer to the first question is.
- (23) ? Roger might know what the answer to the first question is; it depends on whether I got the good answers or not.

This is surprising. Given BASIC CONSTRAINT, there should be no contrast at all between (1)/(4) (on the one hand) and (22)/(23) (on the other). Roger *knows* that if Peggy’s answers are good, then the answer to the first question is b. (Again, he knows what her answer sheet says; he just doesn’t know if it’s right.) He thus knows_G what the answer to the first question is. Thus, the contexts in which natural readings of (1)/(4) arise ought to be exactly those in which

the natural readings of (22)/(23) arise. But the data seem to suggest otherwise. So again: either there is some mysterious feature of the meta-semantics of ‘know’ that explains why G² is a natural constraint question to associate with (1)/(4) but not (22)/(23), or the contrast between the Peggy-centric and Roger-centric ascriptions is evidence that BASIC CONSTRAINT is in need of amendment.

5.2 Knowing and thinking

I suggest we go with the latter option: amend BASIC CONSTRAINT to make it capable of predicting these two contrasts. Out of concerns of space, however, I will not be able to do full justice to the amendment I find most promising. So what follows will have to be taken as a suggestive sketch of the idea, rather than a fully worked out view.

First, a conjecture: holding context fixed, whenever ‘S knows that p’ has a true reading, ‘S thinks that p’ will too. Note: this is not to say that, holding context fixed, whenever ‘S knows that p’ is *felicitous*, ‘S thinks that p’ will be too. If one is in a position to know whether ‘S thinks that p’ is true, then assertions of the form ‘It might be that S thinks that p’ can and often will be infelicitous (because too weak), even when the corresponding ‘It might be that S knows that p’ is just fine. The conjecture is just that if ‘S knows that p’ is *true* at $\langle w, c \rangle$, then so is ‘S thinks that p’.

Second, an observation: the conjecture holds of all the data considered in this paper. Keeping in mind that the claim here is about truth rather than felicity, I leave readers to check this themselves. But here are a few examples of some ‘thinks’ analogs to the knowledge ascriptions we’ve been investigating:

- (24) ✓ If I have the good answers, then I think that the answer to the first question is b. [(1)]
- (25) ✓ One of us thinks that the answer to the first question is b. [(2)]
- (26) ✓ Peggy thinks the answer to the first question is b. [(6)]
- (27) ✓ If Paul is in Paris, then I think he’s in France. [(17a)]

So: supposing the conjecture relating ‘know’ and ‘think’ is true, there would be reason to make it a necessary condition on the truth of ‘S knows that p’ at $\langle w, c \rangle$ that ‘S thinks that p’ be true at $\langle w, c \rangle$ too. This would yield:

THINKING CONSTRAINT ‘S knows that p’ is true at $\langle w, c \rangle$ iff:

- (i) ‘S thinks that p’ is true at $\langle w, c \rangle$.
- (ii) Every world compatible both with S’s *knowledge* (at w) and with the true answer (at w) to the c -determined constraint $\Psi^?$ is a p-world.

And with THINKING CONSTRAINT in place, one could leverage judgments about the falsity of a ‘think’ report into an explanation of the falsity of the corresponding ‘know’ report. This is the strategy I will use to account for the two contrasts.

Applied to the false-seeming (21), (22), and (23) (repeated here)—

- (21) ? Peggy knows who has the good answers.
- (22) ? If I have the good answers, then Roger knows what the answer to the first question is.
- (23) ? Roger might know what the answer to the first question is; it depends on whether I got the good answers or not.

—the idea seems to work. As it turns out, their ‘thinks’ analogs *also* seem false:

- (28) ? Peggy thinks she has the good answers.
- (29) ? If I have the good answers, then Roger thinks the answer to the first question is b.
- (30) ? Roger might think the answer to the first question is b; it depends on whether I got the good answers or not.

And so we seem to have a predictive generalization connecting our judgments about ‘know’ to our judgments about ‘think’. I thus submit that BASIC CONSTRAINT is semantically incomplete. It predicts that there can be true readings of ‘S knows that p’ even when there are no true readings of ‘S thinks that p’. THINKING CONSTRAINT remedies that mistake, and in so doing provides a principled explanation of the unavailability of the good readings of (21)–(23).

5.3 What about ‘think’?

The move from BASIC to THINKING CONSTRAINT raises two salient questions. First, if ‘S knows that p’ entails ‘S thinks that p’, then what is going on with §1’s arguments that our judgments about *History Exam* present *prima facie* counterexamples to principles like BELIEF? Second and relatedly, what can we say about the propositional attitude that is the denotation of ‘think’ on the relevant readings?

I’ll answer the first question in the next section. With respect to the second, essentially what we are asking for is an explanation of *why* there seems to be a true reading of (e.g.) ‘Peggy thinks that the answer to the first question is b’ yet no true reading of ‘Peggy thinks she has the good answers’. And although it’s not clear one *needs* to be able to answer this question to be confident in the judgments or of their use in explaining the two contrasts outlined at the beginning of this section, it would certainly be better if one could. So on that I offer the following.

One’s best guess to a question is the answer one thinks is most likely to be correct. *History Exam* is filled with questions to which its characters can have best guesses. The question ‘What is the answer to the first question?’ admits of four possible answers: a, b, c, and d. Peggy’s

credences in those answers are $\frac{1}{6}$, $\frac{1}{2}$, $\frac{1}{6}$, and $\frac{1}{6}$ respectively. Thus, her best guess to that question is b. By contrast, the question ‘Who has the good answers?’ admits of only two possible answers: Peggy and Pete. Peggy’s credences in each are $\frac{1}{2}$. So it is not the case that her best guess to that question is that she has the good answers, for she has just as much reason to guess that Pete does. Likewise, consider the version of *History Exam* where Roger knows that Peggy’s answer sheet says “1:b” and Pete’s “1:c”, but doesn’t know which is correct. Here Roger’s credences in the four possible answers to ‘What is the answer to the first question?’ are 0, $\frac{1}{2}$, $\frac{1}{2}$, and 0 respectively. It is thus not the case that Roger’s best guess to ‘What is the answer to the first question?’ is b, for he has just as much reason to guess c.

So here is a hypothesis: our judgments about ‘think’ are tracking the underlying facts about best guesses.²⁵ It is because Peggy’s best guess to ‘What is the answer to the first question?’ is b that ‘Peggy thinks that the answer to the first question is b’ is true, and hence why Roger may truly assert (6) (‘Peggy knows what the answer to the first is’). And it is because it is *not* the case that Peggy’s best guess to ‘Who has the good answers?’ is Peggy that ‘Peggy thinks that she has the good answers’ is false, and hence why Roger may *not* truly assert (21) (‘Peggy knows who has the good answers’). Similarly, it is because it is *not* the case that Roger’s best guess to ‘What is the answer to the first question?’ is b that ‘Roger thinks the answer to the first question is b’ is false, and hence why Peggy may *not* truly assert (22) or (23).

I lack the space to further explore the hypothesis that our intuitive judgments about ‘think’ track facts about best guesses. But I submit that connecting ‘know’ to ‘think’ and giving a tentative analysis of the latter in terms of best guesses is a simple, empirically motivated way of helping make sense of some otherwise extremely puzzling data.

6 Epistemology by constraint

6.1 Revisiting the principles

Having explored the semantic and meta-semantic issues raised by the knowledge ascriptions of cases like *History Exam*, I now return to the epistemological principles to which §1 claimed they posed a threat. Since we’ve gone contextualist about knowledge, these are really questions about the truth of the following principles:

‘CREDESCENCE’ For all contexts *c*: If ‘S knows that *p*’ is true in *c*, then ‘S’s credence that *p* is greater than .5’ is true in *c*.

‘BELIEF’ For all contexts *c*: If ‘S knows that *p*’ is true in *c*, then ‘S believes that *p*’ is true in *c*.

‘JUSTIFICATION’ For all contexts *c*: If ‘S knows that *p*’ is true in *c*, then ‘S is in a position to justifiably believe that *p*’ is true in *c*.

²⁵ See Dorst (2019) as well as Holguín (2020) for paper length defenses of the view that there are intimate connections between the facts about thinking and the facts about best guesses.

We'll start with 'BELIEF' as the dialectic surrounding it is more complex than for either 'CRE-DENCE' or 'JUSTIFICATION'. If THINKING CONSTRAINT is true, then 'S knows that p' entails 'S thinks that p'. Whether 'BELIEF' is true then just depends on whether 'S thinks that p' entails 'S believes that p'. And I am more than happy to follow the existing literature in assuming that it does.²⁶ So if THINKING CONSTRAINT is true, then—*pace* the arguments of §1—'BELIEF' is as well.

What then to make of those arguments? Well, there were two of them: first, that *History Exam* shows that 'S knows that p' is sometimes compatible with 'S has arbitrarily low credence that p'; and second, that *History Exam* shows that 'S knows that p' is sometimes compatible with 'S is rationally unwilling to assert or bet that p'. And floating in the background was the thought that 'S believes that p' is incompatible with these things. But this assumption can be rejected. And one who likes THINKING CONSTRAINT should be willing to do so.²⁷

What about 'CRE-DENCE' and 'JUSTIFICATION'? Here I think the decision between MODERATE and THINKING CONSTRAINT is immaterial: given the existence of cases like *History Exam*—cases in which 'S knows that p' expresses a true proposition even when both 'S has greater than .001 credence that p' and 'S is in a position to justifiably believe that p' express false ones—neither principle can plausibly be true in full generality. This should also be unsurprising given the general form of THINKING CONSTRAINT. This theory is designed to accommodate the fact that for just about any proposition p compatible with what S *knows*—i.e., any true proposition p—there will contexts in which 'S knows that p' seems to express a true proposition—give or take the facts about what S's best guesses are. But intuitively neither credence nor justification works like this. It simply does not follow from the fact that p is true (and perhaps that p is S's best guess to some salient question) that there is a true reading of 'S has credence greater than .5 that p' or 'S is in a position to justifiably believe that p'. So absent some decree that ties by fiat 'S knows that p' to 'S has credence greater than .5 that p' and/or 'S is in a position to justifiably believe that p', one should learn to live with the falsity of both 'CRE-DENCE' and 'JUSTIFICATION'. And even if one were to go by fiat, it would not follow that either 'CRE-DENCE' or 'JUSTIFICATION' is true *in the ways epistemologists care about*. The fact that one can conjure up a constraint-based semantics for (e.g.) 'credence' and 'is justified in believing' is no reason to think that the kinds of generalizations that motivated epistemologists to posit principles like CRE-DENCE and JUSTIFICATION are in good standing.

6.2 Revising the principles?

However, even if principles like 'CRE-DENCE' and 'JUSTIFICATION' are false in full generality, there may be narrower interpretations on which they are true. And perhaps it is these interpretations that epistemologists have in mind when they study and defend connections between knowledge, credence, belief, and justification. What are the relevant interpretations? A natural thought is

²⁶ See, e.g., Hawthorne *et al.* (2016), Dorst (2019), Rothschild (2020), Holguín (2020).

²⁷ Cf. Dorst (2019), Holguín (2020).

that they the ones that arise when ‘know’ goes *unconstrained*.²⁸ That is to say: they are principles about **evidence**:

CREDENCE* If S’s **evidence** entails that p, then S’s credence that p is greater than .5.

BELIEF* If S’s **evidence** entails that p, then S believes that p.

JUSTIFICATION* If S’s **evidence** entails that p, then S is in a position to justifiably believe that p.

To the extent one finds this line of thinking plausible, THINKING CONSTRAINT looks epistemologically innocuous. And that is because on this view, one’s **evidence** is one’s *knowledge*—an epistemic relation that is putatively familiar, and that is reasonably believed to satisfy each of the above principles. A proponent of THINKING CONSTRAINT thus has a story in which we properly use the word ‘know’ in all the strange ways discussed so far, but in which those uses tell us little about the worldly relation of concern to epistemologists: *knowledge*.

Of course there would still be some residual worries. For one thing, there remains the question of why we should find the distinction between constrained and unconstrained uses of ‘know’ to be of particular epistemological interest once we’ve convinced ourselves that the full range of readings predicted by THINKING CONSTRAINT are part of the proper use of the concept. For another, one might start to wonder about the epistemology of *knowledge* itself. Much of our thinking about *knowledge* is shaped by thought experiments. For the results of those thought experiments to be probative, it better be that the relevant intuitions involve unconstrained interpretations of ‘know’. But given that unconstrained uses can arise even on unembedded uses (as witnessed by (e.g.) (6)), one wonders about how reliably we can distinguish the unconstrained interpretations from the constrained ones.

But there is also a deeper worry lurking, one that emerges once we place THINKING CONSTRAINT in a larger dialectical context. THINKING CONSTRAINT is somewhat atypical as contextualist theories of knowledge go. Allowing myself to be sloppy with use and mention for the moment: THINKING CONSTRAINT sometimes makes it very *easy* to know things, whereas most of the familiar brands of contextualism are designed to make it sometimes very *hard* to know things.²⁹ What does it mean for knowledge to be easier or more difficult in this way? Well, if we assume *knowledge* corresponds to the “normal” level of difficulty, then we can say that a theory of ‘know’ makes it easy to know things just in case: according to that theory, there are contexts in which ‘S knows that p’ is true even when it’s not the case that S *knows* that p. Likewise, we can say that a theory of ‘know’ makes it hard to know things just in case: according to that theory, there are contexts in which ‘S knows that p’ is false even when S *knows* that p.

²⁸ Alternatively: they are the interpretations that arise when the constraint question is the trivial T?

²⁹ I do not want to oversell the point here. Some epistemic contextualists explicitly try to accommodate “loose” readings—e.g., Schaffer (2005b, 2007); Schaffer & Szabo (2014). While others—e.g., Lewis (1996) and DeRose (2002)—seem to be open to the possibility of extending their contextualist theories to account for such readings.

What is important to observe is that given a body of **evidence**, the mechanisms of constraint can only make knowledge easier to get—never harder. And that’s because the mechanisms of constraint always *shrink* the domain of worlds over which ‘know’ quantifies. So, for instance, if one’s preferred theory of **evidence** is one that ties it to *knowledge*, then on the resulting implementation of THINKING CONSTRAINT, there will never be a context in which S *knows* that p but ‘S knows that p’ is false.

Now given the data considered so far, this is all fine and well: *History Exam*’s knowledge is about as easy as it gets. But the literature on contextualism about knowledge is filled with evidence suggesting that there are contexts in which ‘S knows that p’ is false even though it is uncontroversial that S *knows* that p. Consider the following skeptical puzzle, for instance:

(31) I know that I have hands.

(32) If I know that I have hands, then I know that I am not a handless BIV.

(33) I do not know that I am not a handless BIV.

(31) and (33) both seem to have true readings, while (32) follows from a highly plausible principle like CLOSURE (together with innocuous background assumptions):

CLOSURE If S knows that p and that p entails q, then S is in a position to know that q.

Supposing we want to take both intuitions and CLOSURE seriously, we’re going to have to countenance contexts in which ‘S knows that p’ expresses a false proposition even though S *knows* that p.³⁰ Consequently, if one want one’s contextualist theory of knowledge to predict skeptical uses of ‘know’, one is going to need **evidence** to be something other than *knowledge*—at least in some contexts.

So now a new question arises: is the technical term ‘**evidence**’ context-sensitive or invariant? Supposing it is context sensitive, ‘know’ would be subject to at least two sources of context-sensitivity: context-sensitivity about what is part of one’s **evidence**, and context-sensitivity about what the constraint question is. The question of whether CREDENCE* and JUSTIFICATION* are true would then turn on the question of whether every way of resolving the context-sensitivity of ‘**evidence**’ leaves intact its connections to the other relations.

But supposing instead one gives ‘**evidence**’ an invariantist analysis—which to me seems the more theoretically attractive option—then ‘**evidence**’ will have to context-invariantly denote some extremely weak epistemic relation—something like *factive Cartesian certainty*. By extension, the mechanisms of constraint will have to do all the work in getting ‘know’ to denote a relation like *knowledge*. This would then put the defenders of the **evidence**-centric principles CREDENCE* and JUSTIFICATION* in an awkward position. For although Cartesian certainty plausibly does satisfy each of CREDENCE* and JUSTIFICATION*, the manner in which it does so is less

³⁰ In case it isn’t clear why: If (33) expresses a true proposition in *c*, then given CLOSURE it follows that (31) expresses a false proposition in *c*. But there is no *c* such that ‘I *know* that I have hands’ expresses a false proposition.

than fully satisfying. Now slogans like ‘knowledge requires high credence’, ‘knowledge requires justification’, etc., are true only in virtue of the fact that (e.g.) the relation of *being incapable of rational doubt* requires high credence and justification. I imagine most epistemologists would have hoped for more.

Unfortunately, however, I lack the space to explore these issues further. As such, I will end the discussion of CREDENCE and JUSTIFICATION on a somewhat equivocal note. What is clear is that the ascriptions that have been the focus of this paper show that our ordinary ways of thinking and talking about knowledge are often not responsive to the facts about credence or justification. And that is because giving a proper account of these uses requires a particular kind of contextualism about knowledge—one on which just about any true proposition can count as known given the right setup. It is thus no surprise that these uses of ‘know’ should look so strange from the perspective of conventional epistemology. What remains to be seen is whether there are principled grounds on which the epistemologist can quarantine them.

References

- Blome-Tillmann, Michael. 2009. Knowledge and Presuppositions. *Mind*, **118**(470), 241–294.
- Blumberg, Kyle, & Holguín, Ben. 2019. Embedded Attitudes. *Journal of Semantics*, **36**(3), 377–406. Forthcoming.
- Cohen, Stewart. 1986. Knowledge and Context. *Journal of Philosophy*, **83**(10), 574–583.
- Cohen, Stewart. 1988. How to Be a Fallibilist. *Philosophical Perspectives*, **2**(n/a), 91–123.
- Cohen, Stewart. 1998. Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery. *Australasian Journal of Philosophy*, **76**(2), 289–306.
- Cohen, Stewart. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, **13**(s13), 57–89.
- DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, **52**(4), 913–929.
- DeRose, Keith. 1995. Solving the Skeptical Problem. *Philosophical Review*, **104**(1), 1–52.
- DeRose, Keith. 2002. Assertion, Knowledge, and Context. *Philosophical Review*, **111**(2), 167–203.
- Dorr, Cian, & Hawthorne, John. 2013. Embedding Epistemic Modals. *Mind*, **122**(488), 867–914.
- Dorst, Kevin. 2019. Lockeans Maximize Expected Accuracy. *Mind*, **128**(509), 175–211.
- Dretske, Fred I. 1970. Epistemic Operators. *Journal of Philosophy*, **67**(24), 1007–1023.
- Farkas, Katalin. 2016. Know-Wh Does Not Reduce to Know-That. *American Philosophical Quarterly*, **53**(2), 109–122.
- Farkas, Katalin. 2017. Practical Know-Wh. *Noûs*, **51**(4), 855–870.
- Gettier, Edmund. 1963. Is Justified True Belief Knowledge? *Analysis*, **23**(6), 121–123.
- Gillies, Anthony S. 2010. Iffiness. *Semantics and Pragmatics*, **3**(4), 1–42.
- Ginzburg, Jonathan. 1996. Interrogatives: Questions, Facts and Dialogue. *Pages 385–422 of: Lappin, Shalom (ed), The Handbook of Contemporary Semantic Theory*. Blackwell Reference.
- Greco, John. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Groenendijk, Jeroen, & Stokhof, Martin. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.
- Hamblin, Charles L. 1958. Questions. *Australasian Journal of Philosophy*, **36**(3), 159–168.
- Hawthorne, John. 2000. Implicit Belief and A Priori Knowledge. *Southern Journal of Philosophy*, **38**(S1), 191–210.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, John. 2007. Eavesdroppers and Epistemic Modals. *Philosophical Issues*, **17**(1), 92–101.
- Hawthorne, John, Rothschild, Daniel, & Spectre, Levi. 2016. Belief is Weak. *Philosophical Studies*, **173**(5), 1393–1404.

- Hazlett, Allan. 2010. The Myth of Factive Verbs. *Philosophy and Phenomenological Research*, **80**(3), 497–522.
- Hazlett, Allan. 2012. Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica*, **27**(4), 461–478.
- Heller, Mark. 1995. The Simple Solution to the Problem of Generality. *Noûs*, **29**(4), 501–515.
- Heller, Mark. 1999. The Proper Role for Contextualism in an Anti-Luck Epistemology. *Philosophical Perspectives*, **13**(s13), 115–129.
- Holguín, Ben. 2020. *Thinking, Guessing, and Believing*. Unpublished manuscript.
- Ichikawa, Jonathan. 2011. Quantifiers and Epistemic Contextualism. *Philosophical Studies*, **155**(3), 383–398.
- Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–567.
- Moss, Sarah. 2015. On the Semantics and Pragmatics of Epistemic Vocabulary. *Semantics and Pragmatics*, **5**(8), 1–81.
- Neta, Ram. 2002. S Knows That P. *Noûs*, **36**(4), 663–681.
- Neta, Ram. 2003. Contextualism and the Problem of the External World. *Philosophy and Phenomenological Research*, **66**(1), 1–31.
- Rieber, Steven. 1998. Skepticism and Contrastive Explanation. *Noûs*, **32**(2), 189–204.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, **5**(6), 1–69.
- Rothschild, Daniel. 2020. What It Takes to Believe. *Philosophical Studies*, **177**(5), 1–19.
- Rysiew, Patrick. 2006. Motivating the Relevant Alternatives Approach. *Canadian Journal of Philosophy*, **36**(2), 259–279.
- Rysiew, Patrick. 2016. Epistemic Contextualism. In: Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, winter 2016 edn. Metaphysics Research Lab, Stanford University.
- Schaffer, Jonathan. 2005a. Contrastive Knowledge. *Page 235 of: Gendler, Tamar Szabo, & Hawthorne, John (eds), Oxford Studies in Epistemology 1*. Oxford University Press.
- Schaffer, Jonathan. 2005b. What Shifts? : Thresholds, Standards, or Alternatives? In: Preyer, Gerhard, & Peter, Georg (eds), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.
- Schaffer, Jonathan. 2007. Knowing the Answer. *Philosophy and Phenomenological Research*, **75**(2), 383–403.
- Schaffer, Jonathan, & Szabo, Zoltan Gendler. 2014. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies*, **168**(2), 1–53.
- Schlenker, Philippe. 2004. Context of Thought and Context of Utterance: A Note on Free Indirect Discourse and the Historical Present. *Mind and Language*, **19**(3), 279–304.
- Unger, Peter K. 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- van Elswyk, Peter. 2020. What the Metasemantics of "Know" is Not. *Linguistics and Philosophy*, **43**(1), 69–82.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.