

Knowledge by Constraint

Ben Holguín

June 2018

Abstract

This paper considers some puzzling knowledge ascriptions and argues that they present counterexamples to credence, belief, and justification conditions on knowledge. A novel version of epistemic contextualism is developed to account for these ascriptions, one that takes the interpretation of ‘knows’ to be sensitive to the mechanisms of constraint. Considerations from skeptical knowledge ascriptions and the implicatures of indicative conditionals tell in favor of implementations of the view that confirm the initial impressions about the conditions on knowing.

1 Strange knowledge

Start with a case:

History Exam In preparation for the upcoming multiple-choice final exam, Peggy and Pete have both purchased and subsequently memorized answers sheets from Roger, their corrupt teaching assistant. But moments before the exam Roger shares an unfortunate discovery with both of them: due to his ineptitude, one of Peggy or Pete was given answers to an entirely different exam, and he doesn’t know who it was. (Suppose for the sake of later simplicity that the two answer sheets differ in their answers on every question and that everyone knows this.)

Question 5 is: “In what year did the Berlin Wall fall?” The possible answers are a: 1984, b: 1989, c: 1991, d: 1993. Peggy’s answer sheet says b. Pete’s says c. Peggy is thus the one with the good answers.¹

Putting ourselves in Peggy’s shoes, (1)–(3) all express similar and perfectly natural thoughts:

- (1) If my answer sheet is good, then I know what the answer to question 5 is.

¹ Readers may notice some affinities between *History Exam* and cases discussed by Hawthorne (2000, pp. 202-3; 2004, pp. 68-9). As will be seen shortly, my analysis of *History Exam* is more detailed, and the lessons I draw are quite different.

- (2) One of us knows what the answer to question 5 is. (I hope it's me.)
- (3) I might know what the answer to question 5 is; it depends on whether I got the good answers.

And putting ourselves in Roger (the TA)'s shoes, we might think to ourselves:

- (4) Whoever has the answer sheet that says '5: b' knows what the answer to every question is.

And—still in Roger's shoes—if we discover that Peggy has the good answers, we may even assert something like (5) outright:

- (5) Peggy knows what the answer to question 5 is.

The naturalness of (1)–(5) gives us *prima facie* reason to doubt each of the following widely accepted principles about knowing:²

CREDENCE If S knows that P, then S's credence in P is $> .5$.

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**.³ If Peggy is minimally rational then her credence in the proposition that the answer to question 5 is b can be no higher than .5. She knows that if and only if she got the good answers is the answer to question 5 b, 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 can do no better than guess.

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

MINIMAL CREDENCE If S knows that P, then S's credence in P is $> .001$.

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).

² Readers may notice that each of (1)–(5) invokes knowledge-wh rather than knowledge-that, and that when 'what the answer to question 5 is' is replaced with 'that the answer to question 5 is b', the reports sound noticeably worse. We will have more to say about this in §3.2. For now, however, I will assume that in the context of *History Exam*, if 'Peggy knows what the answer to question 5 is' is true, then so is 'Peggy knows that the answer to question 5 is b'.

³ See, e.g., [Rothschild and Spectre \(2016, 2017\)](#) for arguments that make important use of this principle.

Supposing Peggy is the student with the good answers, (1)–(5) remain perfectly acceptable, even though now Peggy knows that the odds that the answer to question 5 is b are no higher than .001.

Now for BELIEF. If we think that principles like CREDENCE (or at least MINIMAL CREDENCE) hold of belief, then the above considerations provide a direct argument against BELIEF: Peggy knows that the answer to question 5 is b, but her credence in that proposition is not greater than .5 (and could be as low as .001). Belief requires 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 the answer 'b', Peggy lacks any of the ordinary hallmarks of belief. She would not assert that the answer to question 5 is b, nor would she bet on more than even odds that the answer to question 5 is b, nor would she try to convince anyone that the answer to question 5 is b. If Peggy believes the answer to question 5 is b, it is presumably in a sense that is unfamiliar to most of those interested in the truth of BELIEF.⁴

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

The kind of knowledge raised to salience in *History Exam* is a strange one. If we take our intuitions about (1)–(5) at face value, then we have data that presents a strong challenge to a number of widely accepted epistemological principles. Moreover, the case seems not to suggest any straightforward explanation of *why*, supposing the principles are ultimately false, 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 we'd normally call belief. There is no view of knowledge I am familiar with that can account for what is happening in *History Exam*. The aim of this paper is to offer one that does.

The paper is structured as follows. The next two sections expand on the puzzle by offering more cases in a similar vein (§2) and then arguing that certain natural debunking explanations of the data are inadequate (§3). Then we develop a positive theory of knowledge that can account for the data (§4), say a bit about how the background machinery works (§5), and revisit the question of what to think of CREDENCE, BELIEF, and JUSTIFI-

⁴ Some readers have reported the judgment that the 'believes'/'thinks' analogs of (1)–(5) sound fine—i.e., that the sentences still seem true when 'knows' is substituted for 'believes' (or 'thinks') and 'what the answer to question 5 is' is substituted for 'the answers to question 5 is b'. We will return to this issue in a fair bit of detail in §11. There it will be argued that BELIEF may be recoverable after all, but that doing so will require rethinking a number of platitudes about belief.

CATION in light of the view (§6). Next some arguments are presented that challenge the empirical adequacy of the first-pass theory (§7), and then a new, more flexible account is offered in its place (§8). With this view in mind the principles are revisited once again, but this time their prospects are assessed more dimly (§9). Finally we conclude with a discussion of some over-generation worries (§10), an amendment to the view that has the dual effect of insulating the view from the challenge and redeeming BELIEF (§11), and the broader philosophical upshots of the theory of knowledge we are left with (§12).

2 More strange knowledge

2.1 FMSOs

By way of giving the puzzle presented by *History Exam* more bite, it will help to get a sense of how easy it is to imagine cases in which similar patterns of judgment arise. Thought experiments involving perception- and cognition-altering drugs tend to be particularly fertile in this regard:

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. Only Joan’s memory is genuine.

Now consider some thoughts Joan might have to herself (and note the obvious parallels to (1)–(5)):

- (6) If I got the placebo, then I remember what I ate for dinner last night.
- (7) One of us remembers what she ate for dinner last night.
- (8) I might remember what I ate for dinner last night; it depends on whether I got the placebo or the drug.

And taking an experimenter’s perspective, prior to looking at the results:

- (9) Whoever got the placebo remembers what she ate for dinner last night.

And after looking at the results:

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

I will assume that attitude reports involving factive mental stative operators (FMSOs) like ‘remembers that’, ‘sees that’, ‘can hear that’ entail their ‘knows that’ counterparts.⁵ If it is true that S remembers or sees that P, then it is true that S knows that P. Likewise for the analogous ‘-wh’ constructions. Given this assumption, (6)–(10) each entail that Joan knows that she ate fish for dinner last night. It should be clear that Joan’s knowledge, if genuine, is as much a problem for CREDENCE, BELIEF, and JUSTIFICATION as Peggy’s in *History Exam*.

2.2 Implicatures

One more batch of puzzling knowledge ascriptions (for now). Here is a popular and plausible assumption about the implicatures of indicative conditionals and disjunctions, one that I will assume true for the remainder of the paper:

IGNORANCE In uttering ‘If P then Q’ or ‘P or Q’, S typically implicates that S neither knows whether P nor whether Q.⁶

The basic idea behind **IGNORANCE** is simple. If you know whether P (or Q), then why assert the comparatively weak ‘If P, then Q’ or ‘Either P or Q’ when instead you could assert P (or Q) outright? A natural reaction to an utterance of ‘If it is raining in Topeka, then not many people are playing tennis there right now’ is to infer that the speaker is unsure whether it is raining in Topeka, or whether many people are playing tennis there right now. If, for instance, you have good reason to think that the speaker *knows* whether it is raining in Topeka—say, because you saw her look up the weather moments ago—the utterance will tend to seem outright infelicitous.⁷

With that in mind, a case:

Battery One day your laptop suddenly stops working. You take it to a repair store and they tell you that there’s an 80% chance it’s the battery, and that once that’s been replaced everything should be fine. But they also tell you that if the battery doesn’t fix the problem, then the issue could be any number of things and they’ll have to look into it further.

Here are two perfectly appropriate things one might say in light of this information:

⁵ See Williamson (2000, ch. 1).

⁶ See Gazdar (1979) for a canonical discussion of **IGNORANCE**-like principles.

⁷ Utterances made in *dialectical* contexts are exceptions to this generalization. By ‘dialectical contexts’ I mean contexts in which a speaker is rehearsing a bit of reasoning, offering evidence in favor of the propositions she takes herself to know, or hedging to accommodate interlocutors who are insufficiently trusting. These uses are typically of the form: ‘*Even if P*, still Q’, ‘If P, then Q; P, *therefore* Q’, ‘Either P or Q; not P, *therefore* Q’, etc. Whenever I appeal to the principle it will be clear that such contexts aren’t at issue. See Holguín (2018b) for an extended discussion of **IGNORANCE** and its exceptions.

(11) If the new battery doesn't fix the problem, then I don't know what the issue is.

(12) Either the new battery fixes the problem or I don't know what the issue is.

Given IGNORANCE and the fact that (11) and (12) are perfectly natural, it follows that you are ignorant of whether you don't know what the issue with your laptop is. That is to say: (11) and (12) implicate that you do not know whether you do not know what the issue with your laptop is. But this is bizarre, for given standard assumptions about knowledge you are surely in a position to know that you don't know what the issue with your laptop is. After all, you *know* there is a 20% chance the issue is something other than the battery, and that if it isn't the battery it could be any number of things. So how could it be a live possibility that you know what the issue with your laptop is?

Nonetheless (11) and (12) are fine. We have yet more reason to think there are contexts in which subjects count as (possibly) knowing things orthodoxy says they have no business knowing.

3 Two objections

Before getting into the positive account of what is going on with 'knows' in cases like *History Exam*, *Memory Experiment*, and *Battery*, I want to quickly address two objections one might have concerning the probative force of the data.

3.1 First objection: wide-scoping

There is a common first blush reaction to conditional ascriptions like (1) ('If my answer sheet is good, then I know what the answer to question 5 is') (as well as (6) and (11))—and it is that our intuitions about these ascriptions rest on non-obvious subtleties concerning the scope of 'knows' at logical form. 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 question 5 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 question 5 is.

There are at least three reasons to think that the subtleties concerning the scope of attitude verbs in conditionals are beside the point. First reason: 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

ascriptions, (3) is a knowledge ascriptions embedded under a modal, and (4)–(5) 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 entirely non-obvious how the story is supposed to be extended to any of the rest of the data.

Second reason: the wide-scope response founders when the attitude verb is a “thicker” FMSO of perception, such as ‘sees’ or ‘can hear’ (and is already not so great with ‘remembers’). If I’m unsure whether I’ve been given a placebo or a drug that causes auditory hallucinations, I’ll be able to assert or think things like ‘If I got the placebo, then I can hear which album is playing’. But I definitely cannot assert or think things like ‘I can hear whether: if I got the placebo, then this or that album is playing’. The proposition *that if I got the placebo, then this or that album is playing* is just not the kind of thing I can hear, even if it is the kind of thing I can know (and perhaps remember).

Third reason: the wide-scope response also 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. In this 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 question 5 is.
- (1+W) ? Ken and I both know whether: if Pete got bad answers, then the answer to question 5 is b.

Since (1+) is puzzling in the same ways (1) is, and since the former clearly cannot be explained away in terms of a scope ambiguity, wide-scoping does little to help us with the puzzling data.

3.2 Second objection: knowledge-wh vs. knowledge-that

As was observed in footnote 2, each of *History Exam*’s (1)–(5) involves ‘knowledge-wh’ rather than ‘knowledge-that’. The same is true (*mutatis mutandis*) for *Memory Experiment*’s (6)–(10) and *Battery*’s (11)–(12). The reason these examples were stated in terms of knowledge-wh is because the knowledge-that counterparts of (e.g.) (1), (3), and (5) sound markedly worse:⁸

- (1*) (?) If my answer sheet is good, then I know that the answer to question 5 is b.
- (3*) ? I might know that the answer to question 5 is b; it depends on whether I got the good answers.

⁸ (2) and (4) are omitted because neither admits of a straightforward ‘know-that’ translation.

(5*) ? Peggy knows that the answer to question 5 is b.

But the principles for which I claim these data raise problems—*CREDENCE*, *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:

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

What *-WH TO -THAT* says, in short, is that knowing-that is a *necessary* condition on knowing-wh.⁹

I do not think the appropriate reaction to the perceived contrast between the knowledge-wh and knowledge-that ascriptions is to deny *-WH TO -THAT*. Yes, the contrast calls out for an explanation. And unfortunately I lack the space to try and offer one here. But we can be confident *-WH TO -THAT* is true even if we lack an explanation of why it is harder to hear the true readings of the ‘knows-that’ ascriptions than the ‘knows-wh’ ascriptions.

First, if *-WH TO -THAT* were false, then we should expect (13) to have a true uniform interpretation:

(13) ? Although Peggy knows what the answer to question 5 is [namely b], she does not know that the answer to question 5 is b.

But instead it just sounds like a contradiction. I do not know how one would try to explain that on the hypothesis that *-WH TO -THAT* is false.

Second, it is actually not so difficult to hear the true readings of the ‘knows-that’ ascriptions. Three quick points in illustration:

- The ‘knows-that’ analogs of (1)–(5) read naturally when stress is placed on any of ‘Peggy’, ‘question 5’, or ‘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 (14) seems entirely appropriate, while the second seems bizarre:

(14) Q: Does Peggy know what the answer to question 5 is?

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

⁹ With the exception of Farkas (2016a,b), 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 (2016a, p. 3). Note that her putative counterexamples to *-WH TO -THAT* are quite different from the examples of §§1–2.

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

- Finally, the natural answer to ‘Which of the students knows that the answer to question 5 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 question 5 is b, notice that the natural answer to questions like ‘Which of the students is certain the answer to question 5 is b?’ is ‘Neither of them’, not ‘Peggy’.

One who holds the view that the ‘knows-that’ analogs of (1)–(5) are all false must therefore do more than just explain away the appeal of -WH TO -THAT. She 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 at all clear what sort of story can be told here if in fact the ‘knows-that’ ascriptions are context-invariantly false.

A simpler explanation of the asymmetry between the knowledge-wh and knowledge-that ascriptions is that each is most naturally associated with different kinds of resolutions of the context. But in any context in which the knowledge-wh ascription is true, the relevant knowledge-that ascription must be as well.¹⁰ This is the line that will be assumed for the remainder of the paper.

Finally, it is worth stressing that even if -WH TO -THAT were ultimately false, the question of how to give an account of ‘knows-wh’ that captures cases like *History Exam*, *Memory Experiment*, and *Battery* would remain non-trivial. And since it is not at all obvious that there are good reasons to think that the knows-wh analogs of CREDENCE, BELIEF, and JUSTIFICATION are false, arriving at a theory of ‘knows-wh’ that tells us as much would be an interesting result in its own right.

4 Knowledge by constraint

We may now turn to developing an account of ‘knows’ that is capable of explaining what is happening in cases like *History Exam*, *Memory Experiment*, and *Battery*.

I am more or less going to take it for granted that if one wants to vindicate our judgments about sentences like (1)–(12), then one is going to have to be a contextualist about knowl-

¹⁰ An analogy might help here. We more often speak about being sure of things than we do of being certain of them. And that’s because ‘is certain’ seems to be associated with stronger epistemic positions than ‘is sure’. But a view that says that ‘S is sure that P’ can be true while ‘S is certain that P’ is false (or vice-versa) looks rather unpromising. Sentences of the form:

(15) ? I am sure it is raining, but I am not certain it is.

seem only to have false readings. A natural thought is that ‘S is sure that P’ entails ‘S is certain that P’ (and vice-versa), even though the readings we most naturally associate with the two expressions (considered separately) are different.

edge. The alternative is to be something like an *ultra-liberal* invariantist about knowledge, one who thinks that each of BELIEF, MINIMAL CREDENCE, and JUSTIFICATION is false, and that knowledge is in every context as loose as Peggy’s knowledge of the answer to question 5. Aside from its obvious implausibility, the view cannot account for the fact that it is relatively easy to get *bad* readings of just about all of (1)–(12).¹¹ Thus, the important question is what the contextualist machinery that explains the uses of ‘knows’ is supposed to look like.

4.1 The skeleton view

Here is the view I like in abstract.¹² Let us help ourselves to a primitive notion of evidence, which we will call ‘EVIDENCE’. Every agent has a body of EVIDENCE at a given world, which we model as a set of possibilities. To secure the factivity of knowledge, an agent’s EVIDENCE will always include the possibility corresponding to the world the agent is in. In every context, ‘knows’ quantifies universally over a subset of those possibilities, testing to see whether its complement is true throughout the subset. The smaller the subset, the easier it is for every possibility to be consistent with the complement. What changes from context to context is which subset of the agent’s EVIDENCE ‘knows’ quantifies over. In some contexts the subset may be small enough that, by the lights of the context, the agent will know a large number of propositions. But in other contexts the subset might be all of the agent’s EVIDENCE—in which case the subject will (by the lights of the context) know only what is entailed by her EVIDENCE. Intuitively, the readings of ‘knows’ invoked in (1)–(12) are those in which the relevant subset of the initial set of EVIDENCE is small enough to get violations of the principles considered in §1.

What determines which subset of the initial EVIDENCE matters relative to a context? The mechanisms of *non-epistemic constraint*.¹³ Here is how they work. Constraints are polar questions (i.e., yes/no questions). Let $\Psi^?$ denote the question ‘Is it true that Ψ ?’. A set of possibilities Γ is consistent with a constraint $\Psi^?$ at world w just in case: every possibility in Γ agrees with w on the answer to $\Psi^?$. So, if the answer to $\Psi^?$ at w is Ψ , then S ’s EVIDENCE is consistent with the constraint $\Psi^?$ (at w) just in case S ’s EVIDENCE entails Ψ (at w).

We will give a more intuitive characterization of how all this is supposed to work in a moment. But doing so will be easier with the skeleton view of ‘knows’ in hand:

SKELETON ACCOUNT ‘ S knows that P ’ is true at a world w and context c just in case for the c -determined constraint $\Psi^?$, every possibility consistent both with S ’s EVIDENCE and

¹¹ As we just saw, changing the complement of ‘knows’ from a *wh*-clause to the corresponding *that*-clause tends to have this effect.

¹² The view will initially sound much like Lewis (1996)’s. I expect the impression will change. If not, see footnote 18.

¹³ See Dorr and Hawthorne (2013) for a discussion of the mechanisms of non-epistemic constraint in the context of a discussion of epistemic modality. See also Hawthorne (2007) for more on the kind of modal data that is thought to exhibit the constraint phenomenon. (And see also Gillies (2010) and Moss (2015) for further examples of constraint-like accounts of the semantics of epistemic vocabulary.)

with the correct answer to $\Psi^?$ (at the relevant world) is a P possibility.¹⁴

In short, what ‘knows’ tests for is whether its complement (i.e., P) is true throughout the subset of the agent’s EVIDENCE that is consistent whatever proposition is the true answer to the contextually determined constraint question $\Psi^?$.

Here’s a quick example to give a feel for how SKELETON ACCOUNT is supposed to work. Suppose Lexie checked the weather forecast this morning and saw that it said it would rain today. Suppose also that we are understanding the notion of EVIDENCE such that: because Lexie checked the weather forecast and saw that it said rain, every possibility in her EVIDENCE in which the weather forecast is reliable is a possibility in which it is raining. Suppose finally that the weather forecast is in fact reliable. Now consider ‘Lexie knows that it is raining’ in a context in which the constraint (i.e., $\Psi^?$) is ‘Is the weather forecast reliable?’. According to SKELETON ACCOUNT, ‘Lexie knows that it is raining’ is true in this context just in case every EVIDENCE possibility that is consistent with the true answer to ‘Is the weather forecast reliable?’ is a possibility in which it is raining. Since the true answer to ‘Is the weather forecast reliable?’ is ‘Yes’, and since every one of Lexie’s EVIDENCE possibilities in which the forecast is reliable is one in which it is raining, ‘Lexie knows that it is raining’ is true in that context.¹⁵

4.2 EVIDENCE as “ordinary” knowledge

That is how the view works in abstract. But before we fill in the details more and turn to (1)–(12), I want to first put a candidate theory of EVIDENCE on the table. Doing so will give us clearer intuitions about the view’s predictions.

Let us suppose that one’s EVIDENCE is one’s *ordinary knowledge*. I will use ‘*knows*’ (that is, the *asterisks*) to denote ordinary knowledge and related notions (e.g., *epistemic possibility*, *ignorance*, etc.). *Ordinary knowledge* is a black box, though it is less of a black box than EVIDENCE. The basic idea is that it’s just what your typical moderate invariantist about knowledge thinks knowledge is. But the following stipulations about *knowledge* should help pin down the notion a bit more. I *know* that I have hands, who the current president of the United States is, where I’ll be living in a week, and so on. Gettier cases are not cases of *knowledge*. I don’t *know* that my lottery ticket will lose (even if in fact it will). Vagueness precludes *knowledge*. *Knowledge* satisfies CREDENCE, BELIEF,

¹⁴ The relevant world will in most cases be w , the world of evaluation. It is only when ‘knows’ is embedded in something like a conditional or an epistemic modal that the relevant world is sometimes shifted.

¹⁵ SKELETON ACCOUNT finds semantic analogies between ‘knows’ and other expressions that are subject to the mechanisms of constraint. If [Dorr and Hawthorne \(2013\)](#)’s claims are right (in broad outline), then epistemic modals like ‘must’ and ‘certainly’ are the closest semantic analogs to ‘knows’. This makes ‘knows’ a kind of adverbial quantifier (like ‘always’ or ‘must’) rather than a domain quantifier (like ‘all’ and ‘every’—cf. [\(Ichikawa, 2011\)](#)). See [Schaffer and Szabo \(2013, §§2.4, 3.1, and 4.5 in particular\)](#) for discussion of the advantages modeling ‘knows’ along the lines of a quantificational adverb (as SKELETON ACCOUNT does), particularly as concerns standard semantic challenges to epistemic contextualism.

and JUSTIFICATION. Thus—though also stipulatively, if need be—Peggy does not *know* what the answer to question 5 is, Joan does not *know* or *remember* what she ate for dinner last night, and the ‘knows’ of (11)–(12) does not denote *knowledge*.

EVIDENCE as *knowledge* gives us the following slightly fleshier implementation of SKELETON ACCOUNT:

MODERATE CONSTRAINT ‘S knows that P’ is true at a world w and context c just in case for the c -determined constraint $\Psi^?$, every possibility consistent both with what S *knows* and with the correct answer to $\Psi^?$ (at the relevant world) is a P possibility.

The way MODERATE CONSTRAINT works is simple. If you want to know whether ‘S knows that P’ is true, look at all the possibilities compatible with what S *knows*, prune off those that are inconsistent with the true answer to the contextually determined $\Psi^?$, and see if any remaining possibilities are $\neg P$ possibilities. S knows that P if and only if there are no such possibilities. A more illustrative but less precise way of thinking about it: if ‘S knows that P’ is true when the contextually 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.¹⁶

4.3 Accounting for the data

We now turn to explaining how a view like MODERATE CONSTRAINT can account for the data of §§1–2. For now we will help ourselves to assumptions about what the contextually relevant $\Psi^?$ is on the natural readings of the various ascriptions. The next section will consider whether those assumptions can be argued for on independent grounds.

Beginning with *History Exam*, let us suppose the natural readings of (1)–(5) are those in which the constraint $\Psi^?$ is along the lines of ‘Does Peggy have good answers?’. Peggy has the good answers, so ‘Peggy knows what the answer to question 5 is’ is true (on its constrained reading) just in case: all the *epistemic possibilities* (for Peggy) that agree with the true answer to the question ‘Does Peggy have good answers?’ are *epistemic possibilities* in which the answer to question 5 is b. This is indeed the case. Peggy *knows* that if the answer to question 5 is b, then her answers are good.¹⁷ Thus, it is not an *epistemic possibility* for Peggy that she has the good answers and yet the answer to question 5 is something other than b. Since Peggy in fact has the good answers, all the *epistemic possibilities* consistent with the true answer to the constraint question are *possibilities* in which the answer to question 5 is b. Thus, ‘Peggy knows what the answer to question 5 is’ is true in context, as desired. This gives us (4) and (5) straight away. (1)–(3) embed ‘Peggy knows what the

¹⁶ This way of thinking about things misleads in cases where the satisfaction of the counterfactual’s antecedent has downstream consequences for the rest of S’s *knowledge*. For example: maybe the nearest world in which S *learns* the answer to $\Psi^?$ is one in which she becomes an amnesiac.

¹⁷ Recall that Peggy *knows* from Roger that if her answer to question 5 is correct, then her answers to every other question are correct too.

answer to question 5 is' under a conditional, quantifier, and modal respectively, and it is straightforward to see that the ascription remains true on these embeddings when it takes the relevant constrained reading. The same applies (*mutatis mutandis*) to the natural readings of (6)–(10) in the context of *Memory Experiment*, but with a contextually determined constraint along the lines of 'Did Joan get the drug?'

Now consider *Battery*'s puzzling sentences:

(11) If the new battery doesn't fix the problem, then I don't know what the issue is.

(12) Either the new battery fixes the problem or I don't know what the issue is.

Recall that you have been told that there is an 80% chance the battery is the cause of your laptop's malfunctioning, but also that if it isn't the battery it could be any number of things. By IGNORANCE, indicative conditionals and disjunctions carry the implicature that you do not know their consequents or disjuncts. But given your *evidence*, you are in a position to *know* that you don't *know* what is wrong with your laptop. Since neither (11) nor (12) seems to be in violation of IGNORANCE, we have reason to believe that neither sentence's 'knows' denotes *knowledge*.

Enter MODERATE CONSTRAINT. Suppose the 'knows' relevant in the assessment of IGNORANCE is *knowledge* and that the constraint on the 'knows' that appears in the consequent of each of (11) and (12) is something along the lines of 'Is the battery the source of the problem?'. On this reading, 'I don't know what the issue is' is true just in case: the *epistemic possibilities* consistent with the true answer to the question 'Is the battery the source of the problem?' do not all agree on what the issue with the laptop is. Whether that condition holds depends on whether the battery is the source of the problem or not. If the battery *isn't* the source of the problem, then the *epistemic possibilities* consistent with the true answer to the constraint question disagree on what the issue is—in which case the condition holds. But if the battery *is* the issue, then trivially the *epistemic possibilities* consistent with the answer to the constraint question agree that the issue is the battery—in which case the condition doesn't hold. Since you don't *know* whether the battery is the source of the problem, you don't *know* whether the condition holds. Therefore, you don't *know* whether 'I don't know what the issue is' is true. MODERATE CONSTRAINT thus predicts that (11) and (12) carry the desired implicatures.

MODERATE CONSTRAINT provides a simple and intuitive explanation of the behavior of 'knows' in (1)–(12). The core idea is that 'knows' is susceptible to constrained readings. When it takes such a reading, the relevant propositions need only be true at contextually determined subsets of *epistemic space*, rather than all of it. This in turn explains why subjects like *History Exam*'s Peggy and *Memory Experiment*'s Joan can be ascribed knowledge of propositions for which they possess at best dubious evidence, and why in the context of cases like *Battery* we can speak and think as if we are *uncertain* about what we know

even when we are *not* *uncertain* about what we *know*. Absent evidence that existing theories of knowledge are similarly well-situated, we already have reason to prefer some variant of SKELETON ACCOUNT to its competitors.¹⁸

Still, two questions naturally suggest themselves at this stage of the discussion: first, why are constraint questions, well, *questions*? Why couldn't we just use propositions? And second, do we have any story about why the constraint gets the value it does in each of the cases? The next section answers these two questions in turn. Afterwards we will return to the question of what to think of principles like CREDENCE, JUSTIFICATION, and BELIEF in light of MODERATE CONSTRAINT.

5 Questions and concerns

5.1 Why constraint *questions*?

The need for a question-like restriction mechanism is made most vivid by sentences like *Battery's* (11) and (12). If context supplied a proposition to restrict the relevant body of *evidence*—say, the proposition that the battery *isn't* the problem—then you would be in a position to *know* that you don't know (in the relevant restricted sense) what the issue with your laptop is. But again, the reason (11) and (12) are not in violation of IGNORANCE is that you *don't* *know* whether you don't know what the issue is.

The same point can be made more sharply without recourse to a principle like IGNORANCE. Consider Peggy's utterance of (16) in the context of *History Exam*:

(16) There's a 50% chance I know what the answer to question 5 is.

Like, (1)–(5), (16) reads perfectly naturally in context, and so we want our theory of 'knows' to be able to account for it. If (16)'s 'knows' denoted *knowledge*, then it would be false: by hypothesis, there is a 100% chance Peggy doesn't *know* what the answer to question 5 is. So it must take a constrained reading. However, if the constraint were generated by a *proposition*—say, the proposition that Peggy has the good answers—then (16) would continue to remain false. And that's because there is a 100% chance (not 50%) that all the

¹⁸ How does MODERATE CONSTRAINT compare to extant forms of epistemic contextualism, e.g.: the views of Cohen (1986, 1988, 1999); DeRose (1992, 1995); Lewis (1996); Neta (2003); Schaffer (2007); Blome-Tillmann (2009); Ichikawa (2011); Schaffer and Szabo (2013)? The short answer is this: none of these theorists use the mechanism of constraint to generate the different resolutions of 'knows', and plausibly only views along the lines of Lewis's could in principle. For reasons I lack the space to get into fully, it seems clear that (1)–(12) are amenable neither to a "shifting standards" model of epistemic contextualism (as Cohen, DeRose, and Neta each favor) nor to a "shifting question under discussion" model of epistemic contextualism (as Schaffer (and Szabo) favors). Lewisians (e.g. Lewis, Blome-Tillman, and Ichikawa) can in principle make use of the mechanism of constraint, though that is mostly just because the kernel of the Lewisian approach to epistemic contextualism is the very general thought that 'knows' is a quantifier over a domain of possibilities. What is important is that no existing Lewisian uses a question-like mechanism (such as constraint) to generate the various resolutions of 'knows', one whose content is sensitive to the linguistic environment in which the relevant knowledge ascription is embedded. The next section will explain the importance of this feature of the view in more detail.

possibilities consistent both with what Peggy *knows* and with the proposition that Peggy has the good answers are possibilities in which the answer to question 5 is b. To get it to be that there is only a 50% chance that the proposition expressed by ‘I know what the answer to question 5 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 question 5 is b. And that is exactly what we achieve by using questions. Peggy is 50% certain that the answer to ‘Does Peggy have the good answers?’ is the proposition that Peggy has the good answers. She is thus 50% certain that all her *epistemic possibilities* consistent with the true answer to that question are *possibilities* in which the answer to question 5 is b.

More generally, it is no coincidence that all of the first-personal knowledge ascriptions among (1)–(12) are embedded in constructions typically used to express uncertainties: conditionals, disjunctions, epistemic modals, etc. And that is because in each case the speaker is uncertain of whether they stand in the contextually determined epistemic relation to the proposition expressed by the knowledge ascription’s complement. It is this fact that requires that constraints have the structure of questions rather than propositions. If constraints had the structure of propositions, there would be no room for uncertainty about whether the agent stands in the contextually determined epistemic relation to the proposition expressed by the complement of the embedded knowledge ascription.

5.2 The meta-semantics of constraint

We are now in a position to look more closely at some of the assumptions that went into MODERATE CONSTRAINT’s treatment of (1)–(12). In particular, we can now ask what explains why constraint questions get the values they do in context. §4’s story about how MODERATE CONSTRAINT accounts for the data of §§1–2 made substantive assumptions on this issue. In particular, it assumed that the natural readings of the *History Exam* ascriptions are those in which the constraint is something like ‘Does Peggy have the good answers?’, that for *Memory Experiment* it’s something like ‘Did Joan get the drug?’, and that for *Battery* it’s something like ‘Is the battery the source of the problem?’. It is natural to wonder what, if anything, explains why these are the constraint questions associated with the natural readings of (1)–(12).

An immediate answer draws on basic abductive considerations: in each case the hypothesis explains the data, and does so in a way that is faithful to the underlying semantic phenomenology. Absent a plausible competitor, we should settle for what works—and I know of no better treatment of the data. But I won’t pretend that’s the same thing as having a genuine explanation. Can we improve our answer?

A natural thought is to try and derive a theory of the meta-semantics of constraint from more familiar theories of contexts and context-sensitivity. Perhaps the contextually determined $\Psi^?$ is always just the contextually salient “question under discussion”, in the style

of (e.g.) Roberts (2012). Supposing the workings of constraint could be reduced to the workings of QUDs, then so long as the constraint questions posited in §4’s explanation of (1)–(12) could plausibly be the relevant QUDs associated with those contexts, MODERATE CONSTRAINT’s account of the data would look considerably less ad hoc.

Unfortunately, there are good reasons to believe that the workings of constraint cannot generally be reduced to the workings of QUDs. To see why we will focus on *History Exam*. For the story to work in this case, it would have to be that the QUD associated with the natural readings of (1)–(5) is ‘Does Peggy have the good answers?’ And maybe that’s right. But issues arise when we consider slight variants on the original ascriptions. Compare (1) to its “tautologous” analog (1’):

- (1) ✓ If I have the good answers, then I know what the answer to question 5 is.
- (1’) ? If I have the good answers, then I know whether I have the good answers.

(1’) is significantly worse than (1). The issue is that if the contextually determined constraint were always just the salient question under discussion, we would not expect to see this contrast. Here is why. Given MODERATE CONSTRAINT and the details of *History Exam*, (1) is true whenever the contextually determined constraint is ‘Does Peggy have the good answers?’.¹⁹ This means that the natural contexts associated with (1) are those in which (1’) is true as well. So the difference in our judgments about (1) and (1’) must therefore be due to the differences in the kinds of contexts we associate with their natural readings. And therein lies the problem for the constraints-as-QUDs model. By hypothesis, the QUD associated with the natural readings of (1) is ‘Does Peggy have the good answers?’. To account for the fact that (1’) sounds false, it would have to be that the QUD associated with its natural readings *can’t* be ‘Does Peggy have the good answers?’. But this is not at all what we should expect: the most straightforward difference between (1) and (1’) is that (1’) makes salient the question of who has the good answers, while (1), by contrast, merely makes salient the question of who knows the answer to question 5. That is to say: if one of (1) and (1’) were more likely to be associated with the QUD of ‘Does Peggy have the good answers?’ it would be (1’)—exactly the opposite of what is suggested by the data. We thus have some reason to expect a question to be less likely to be used as a constraint when it is simultaneously the salient question under discussion.²⁰

In fact, the inverse correlation between a question’s working as a constraint and its being a QUD may be what *explains* the general badness of “tautologous” constructions of the form:

¹⁹ The argument is trivial. If Peggy has the good answers and the constraint question is ‘Does Peggy have the good answers?’, then obviously every *epistemic possibility* compatible with the true answer to that question is one in which Peggy has the good answers.

²⁰ Perhaps a more flexible understanding of QUDs would allow one to get around this problem. The worry would then be that the resulting notion is only superficially illuminating as regards the meta-semantics of constraint. But absent concrete proposals it is difficult to assess the possibilities.

(17) ? If P, then I know whether P.²¹

The rough thought here is that when the QUD is $P^?$ (or something sufficiently close to it), then knowledge ascriptions that are constrained by the true answer to that question are likely to seem trivial. Hence the preference against resolving the context that way.

Anyway, without avail of QUDs I am not optimistic about the general strategy of repurposing more familiar contextualist machinery in the pursuit a meta-semantic theory of the workings of constraint. Still, I do not think all hope is lost. I believe there are patterns in the data strong enough to warrant some predictive generalizations. In particular—though with two caveats—the following looks to be a fairly robust generalization about the resolutions of constraint in context. It is that the natural readings of sentences of the form

(20) If P, then S knows whether Q.

(21) Either $\neg P$, or S knows whether Q.

(22) It might be that S knows whether Q; depends on whether P.

tend to be those in which the contextually determined constraint is $P^?$. That is to say, when ‘knows’ is embedded in certain linguistic environments—conditionals, disjunctions, and certain uses of epistemic modals (those with continuations that make a question salient)—the content of the constraint question tends to be supplied by other parts of the environment: the conditional’s antecedent, the negation of the other disjunct, whatever follows the ‘depends on whether...’ construction, etc. This is exactly what we see with *History Exam*, *Memory Experiment*, and *Battery*. Examples may be multiplied with ease:

(23) If Trump was born in the 1930s, then we know whether he is lying about his age.

²¹ Mysteries abound with sentences like (17). For one thing, ever so slight variations in the embedded ascription’s complement have the effect of noticeably improving the conditional. Change the complement to P’s negation, or to a disjunction of the form ‘P or ($\neg P$ and Q)’, and the sentence dramatically improves:

- (17) a. ? If it rained yesterday, then I know whether it rained yesterday.
b. ✓ If it rained yesterday, then I know whether it was sunny yesterday.
c. ✓ If it rained yesterday, then I know whether it either rained or snowed yesterday.

Why should ‘If P, then I know whether P’ be bad while ‘If P, then I know whether not-P’ and ‘If P, then I know whether either P or Q’ are both fine? Notice too that you see basically the same pattern of data with ‘must’ and ‘certainly’:

- (18) a. ? If it rained yesterday, then it must have rained yesterday.
b. ✓ If it rained yesterday, then it must not have been sunny yesterday.
(19) a. ? If it rained yesterday, then it certainly rained yesterday.
b. ✓ If it rained yesterday, then it certainly either rained or snowed yesterday.

What is suggested by these observations is that indicative conditionals are generally less likely to exhibit their normal restriction-like effects on an expression Φ when they take the tautologous’ form ‘If P, Φ P’.

- (24) Either very improbably we are the only intelligent beings in the universe, or we know that we are not alone.
- (25) We might know where to find our missing cat; depends on whether she escaped through the front door or the back window.

Given that the generalization is robust, it is reasonable to appeal to it in account for the data in the way MODERATE CONSTRAINT does. Thus, with two caveats to be discussed momentarily, it is a prediction of SKELETON ACCOUNT that sentences of the form of (20)–(22) will have true readings whenever it is part of S’s EVIDENCE that P entails Q.

Here are the two caveats. First, the generalization fails for tautologous conditional knowledge ascriptions along the lines of (17) (‘If P, then I know whether P’).²² Second, although the generalization will more or less invariably hold when S, the subject of the ascription, also happens to be the speaker, there will be many counterexamples when the speaker and subject are distinct. §10 will discuss this worry in more detail, and §11 will offer an amendment to SKELETON ACCOUNT to account for it.

For present purposes neither caveat is of concern. None of (1)–(12) is a tautologous conditional knowledge ascription, and all the ascriptions that fit the pattern of (20)–(22) have the speaker coinciding with the subject. Thus, with regards to the question of why the constraint associated with the natural readings of (1)–(5) in the context of *History Exam* is ‘Does Peggy have the good answers?’, the answer is that most of those sentences fit the pattern of (20)–(22).

But what about those that don’t fit the pattern? The generalization does nothing to explain why the *unembedded* (5) (‘Peggy knows what the answer to question 5 is’) seems true when uttered by Roger, an onlooker who *knows* that Peggy has the good answers. What sentences like (5) show is that sometimes the content of the constraint has to be supplied by features of the extra-linguistic context. There is lots of vague stuff to be said about the circumstances in which we might expect these sorts of ascriptions to sound more natural. For instance: ascriptions like (5) tend to sound better when the ascriber’s interest is in how the subject will *behave*—e.g., whether Peggy will write ‘b’ for question 5—rather than with the quality of the subject’s *evidence*—e.g., how good Peggy’s reasons are for believing that the answer to question 5 is b. Unfortunately, however, I doubt we can do much better than that. The data just seems to be too unruly to be amenable to mechanical explanations in the style of the one offered for (20)–(22).

So this is where we are on the issue of the meta-semantics of constraint. Setting aside the two caveats from above, we have some general rules for predicting which constraint question will be operative in context *for knowledge ascriptions embedded in conditionals, disjunctions, and modals*. And for the rest of the ascriptions we have vague heuristics and just-so stories that appeal to salient features of the circumstances in which the ascriptions

²² See footnote 21.

occur. One might take this to be evidence that MODERATE CONSTRAINT is problematically incomplete. But against rushing to this hasty conclusion it is worth stressing two points.

First, *no one* has a tidy story that explains why (e.g.) (5) should seem true given the facts of *History Exam*. Peggy lacks most all of the features philosophers associate with knowledge. So it is no strike against MODERATE CONSTRAINT in particular that it must appeal to extra-linguistic features of context to account for it. And unlike most other views, MODERATE CONSTRAINT can at least accommodate the data.

Second, one ought not think that a contextualist semantics for an expression ought to be accepted only if its meta-semantics has already been worked out. That is to say, we need not have 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 in some room—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 ‘knows’.²³

6 The principles, revisited

With a working theory of ‘knows’ that can accommodate (1)–(12), as well as some tentative generalizations about the value of the constraint parameter in context, we may now return to the question of what to make of the principles that §§1-2 took to be threatened by cases like *History Exam*, *Memory Experiment*, and *Battery*:

CREDENCE If S knows that P, then S’s credence in P is $> .5$.

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.

If these principles are true if and only if they are true on *every* resolution of ‘knows’, then they are either false or of a highly unfamiliar form. The reason why is that the following is a general property of SKELETON ACCOUNT: propositions that are true at arbitrarily small regions of one’s EVIDENCE can count as knowledge given the right constraint question.²⁴ But

²³ Ichikawa (2011) makes a similar point in defending Lewis’s (1996) contextualist semantics for ‘knows’ from common objections to the meta-semantic Lewis happened to pair it with.

²⁴ We will return to the obvious over-generation concerns raised by this fact in §10.

intuitively none of credence, belief, or justification works like this. So absent some decree that ties these notions to knowledge by fiat, we simply have no reason to expect that because P is true, there is thereby a sense in which I have credence greater than .5 that P, or a sense in which I am justified in believing that P.²⁵

But what about a less ambitious reading of the principles? Perhaps when epistemologists study and defend principles like BELIEF or JUSTIFICATION, what they have in mind is a limited (but privileged) range of interpretations of ‘knows’. So long as the principles are true on these interpretations, the thought goes, they are true in the way that “matters”. On the natural ways of developing the thought, the privileged resolutions of ‘knows’ are presumably the *unconstrained* resolutions. That is, the resolutions on which ‘knows’ quantifies over the whole of one’s EVIDENCE.

To the extent we find this line of thinking plausible, MODERATE CONSTRAINT looks to be friendlier to an invariantist epistemology than we might have initially expected. And that is because on MODERATE CONSTRAINT one’s EVIDENCE is one’s *knowledge*—an epistemic relation that is putatively familiar, and that we know satisfies each of the above principles. A proponent of MODERATE CONSTRAINT thus has a story in which we properly use the word ‘knows’ in all the strange ways witnessed in §§1–2, but in which those uses tell us little about the worldly relation of concern to epistemologists: *knowledge*.

Of course, there remains the question of why we should find the distinction between constrained and unconstrained uses to be of particular epistemological interest once we’ve convinced ourselves that the full range of readings predicted by MODERATE CONSTRAINT are part of the proper use of the verb ‘knows’. But rather than address this question directly, I will instead argue that MODERATE CONSTRAINT is not the correct implementation of SKELETON ACCOUNT. There are strong reasons to prefer filling in the EVIDENCE box with an epistemic relation that is *stricter* than *knowledge*. I will then argue that once we’ve adopted the proper conception of SKELETON ACCOUNT, the strategy outlined here for recovering the epistemological interest of principles like BELIEF and JUSTIFICATION is unpromising.

7 Against MODERATE implementations of SKELETON ACCOUNT

MODERATE CONSTRAINT is somewhat atypical as contextualist theories of knowledge go. Allowing ourselves to be sloppy with use and mention for the moment: MODERATE CONSTRAINT sometimes makes it very *easy* to know things, whereas most of the familiar brands

²⁵ Or at least it is not clear that we have reason to expect these things given that ‘credence’, ‘justification’, etc. are being used *in the ways epistemologists care about*. Again: the fact that one can conjure up a constraint-based semantics for (e.g.) ‘believes’ and ‘is justified in believing’ is no reason to think that the kinds of generalizations that motivated epistemologists to posit principles like BELIEF and JUSTIFICATION carry through to the full range of the uses of ‘knows’.

of contextualism are designed to make it sometimes very *difficult* to know things.²⁶ 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 ‘knows’ 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 ‘knows’ 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.

What is important to observe is that given an EVIDENTIAL base, 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 possibilities over which ‘knows’ quantifies. So, for instance, if our preferred theory of EVIDENCE is one that ties it to *knowledge*, then on the resulting implementation of SKELETON ACCOUNT (namely MODERATE 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—the ascriptions and implicatures in *History Exam*, *Memory Experiment*, and *Battery*—this is all fine and well: that kind of knowledge is about as easy as it gets. But there is strong evidence that sometimes it is harder to know things than it is to *know* them. Less sloppily: there is strong evidence that there are many contexts in which intuitively ‘S knows that P’ is false even though it is uncontroversial that S *knows* that P—a fact that MODERATE CONSTRAINT cannot account for. Supposing we want to take such intuitions seriously, it follows that EVIDENCE needs to include a greater variety of possibilities than what is afforded by *knowledge*.

This section will present two bits of data that suggest that the epistemic relations denoted by ‘knows’ are sometimes more demanding than *knowledge*. The first is familiar from the existing literature on epistemic contextualism. The second is novel and draws on considerations from principles like §2.2’s IGNORANCE.

7.1 Skeptical puzzles

Consider the following two familiar skeptical puzzles:²⁷

- (26) I know that I have hands.
- (27) If I know that I have hands, then I know that I am not a handless BIV.
- (28) I do not know that I am not a handless BIV.

and:

²⁶ I do not want to oversell the point here. Some epistemic contextualists explicitly try to accommodate “loose” readings (e.g. Schaffer (2005b, 2007); Schaffer and Szabo (2013)). While others (e.g. Lewis (1996); DeRose (2002)) are explicitly open to the possibility of extending their contextualist theories to account for such readings.

²⁷ See, e.g., Klein (2015) for an introduction to the vast literature on skepticism.

- (29) I know that Donald Trump is the president of the United States.
- (30) If I know that Donald Trump is the president of the United States, then I know that he has not had a fatal heart attack in the past few moments.
- (31) I do not know that Donald Trump has not had a fatal heart attack in the past few moments.

In each case the first and third members of the triad are intuitively true, while the second 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.²⁸

Yet the three members of the triad are jointly inconsistent, and so something has to give.

In light of the puzzle we have four options: (i)-(iii): deny the first, second, or third members of the triad respectively; or (iv): go contextualist about ‘knows’.

Skeptical invariantists about ‘knows’ go for option (i) and deny that either (26) or (29) is true. Moderate invariantists go for option (iii) and deny that either (28) or (31) is true. Foes of CLOSURE go for option (ii) and allow for both the first and third members of the triads to be true in the same context, but deny that the second members are. Finally, epistemic contextualists (of the kind that are of interest) go for option (iv) and take CLOSURE (and thus (27) and (30)) to be true in every context, but deny that the first and third members are ever true in the same context. To explain the recalcitrant intuitions, the contextualist also claims that the natural readings of the first and third members of the triads occur in different contexts.

Before seeing where SKELETON ACCOUNT and its MODERATE implementation fits into this picture, let us introduce some new notation. Let ‘knows_{Ψ?}’ denote what SKELETON ACCOUNT says ‘knows’ denotes when the contextually operative constraint is Ψ[?]. So, ‘S knows_{Ψ?} that P’ is true just in case all the possibilities that are consistent with S’s EVIDENCE and that agree with the relevant world on the answer to Ψ[?] are P worlds. Thus, given MODERATE CONSTRAINT, ‘knows_{Ψ?}’ denotes *knowledge* constrained by the answer to Ψ[?].

MODERATE CONSTRAINT falls into the camp of those who take option (iii) and deny that (28) or (31) is ever true. CLOSURE comes for free: no matter the Ψ[?], if S knows_{Ψ?} that P and that P entails Q, then S knows_{Ψ?} that Q.²⁹ And since *knowledge* is supposed to reflect our moderately invariant intuitions about knowledge, it follows that a proponent of MODERATE CONSTRAINT must think that (28) and (31) are false in *every* context.

But the intuitions that (28) and (31) are true are quite strong. The fact that a proponent of MODERATE CONSTRAINT cannot accommodate them is non-trivial strike against the

²⁸ For an introduction to some of the many issues involved with CLOSURE, see, e.g., Hawthorne (2005).

²⁹ A bit more detail: If S knows_{Ψ?} that P, then every *epistemic possibility* consistent with the answer to Ψ[?] is a P possibility. If S knows_{Ψ?} that P entails Q, then every P *epistemic possibility* consistent with the answer to Ψ[?] is also a Q possibility. Thus, if S knows_{Ψ?} that P and that P entails Q, S must also know_{Ψ?} that Q.

theory. The point is especially pressing given that MODERATE CONSTRAINT is already a form of contextualism about ‘knows’. It is not like the losses incurred by denying the intuitive judgments about (28) and (31) are offset by gains in avoiding whatever other problems contextualism is typically saddled with.³⁰

7.2 Indicative conditionals and IGNORANCE

The next consideration in favor of a more flexible implementation of SKELETON ACCOUNT concerns our intuitions about indicative conditionals. Recall IGNORANCE from the discussion of *Battery* in §2.2:

IGNORANCE In uttering ‘If P then Q’ or ‘P or Q’, S typically implicates that S neither knows whether P nor whether Q.

With IGNORANCE in mind, consider the following ordinary indicative conditionals:

- (32) If Oswald didn’t kill Kennedy, then someone else did.
- (33) If Trump is still the president of the United States, then he hasn’t had a fatal heart attack in the past few moments.

—and also some more philosophically loaded ones:

- (34) If nature isn’t uniform, then induction is not a reliable means of forming true beliefs.
- (35) If you are in the good case, then you know whether you have hands.

Each is perfectly acceptable in a number of ordinary contexts. Thus, given IGNORANCE, the contexts in which one felicitously asserts any of (32)–(35) are contexts in which one is ignorant of the truth-value of both their antecedents and consequents. But by stipulation you *know* that Oswald killed Kennedy, Trump is president, nature is uniform, and that you are in the good case. It follows that IGNORANCE’s ‘knows’ must sometimes denote a stricter epistemic relation than *knowledge*. And this means that anyone who accepts IGNORANCE should take (32)–(35) to show that MODERATE CONSTRAINT supplies an insufficiently wide range of resolutions of ‘knows’.³¹

³⁰ See, for example, the charge raised by (e.g.) Schiffer (1996); Hawthorne (2004); Williamson (2005) that contextualism about ‘knows’ implies that the phenomenon of semantic blindness is much more widespread than is antecedently plausible. I lack the space to deal with such objections here (though see Schaffer and Szabo (2013) and Greenough and Kindermann (2017) for some help on that front). For present purposes the important point is just that the friend of MODERATE CONSTRAINT already has to answer these sorts of worries, and so doesn’t lose much by expanding the scope of her theory to cover ordinary intuitions about skeptical puzzles.

³¹ Sentences like (32)–(35) might inspire one to have doubts about IGNORANCE along the following lines. If I know that Jane is usually at home at this hour, but that if she isn’t at home then she’s either at work or on the tennis courts, I can felicitously assert something like:

(J) If Jane isn’t at home, then she might be at work.

8 From MODERATE to SKEPTICAL CONSTRAINT

8.1 EVIDENCE as Cartesian knowledge

We have reason to look for a different implementation of SKELETON ACCOUNT'S EVIDENCE relation, one that allows for resolutions of 'knows' on which subjects fail to know what they *ordinarily know*. Here is the new candidate: one's EVIDENCE is one's *Cartesian knowledge*. I will use '**knows**' (that is, **boldface**) to denote Cartesian knowledge and related notions (e.g., **epistemic possibility**, **ignorance**, etc.) Like *ordinary knowledge*, **Cartesian knowledge** is something of a black box. But here are some of the roles that will help characterize it. **Knowledge** is what the epistemological skeptic thinks knowledge is. Expressions associated with it include 'What is known *for sure*', 'What is known *without a doubt*', and so on. You don't **know** that you have hands, who the president of the United States is, or even whether there is an external world. With the exception of the proposition that you exist and some propositions about your conscious experience, you are **ignorant** of just about every contingent proposition. Candidates for **knowledge** are facts about your conscious experience and whatever can be "known" a priori—conceptual truths, facts about logic and math, etc.³² **Knowledge** thus (somewhat cheaply) satisfies CREDENCE, BELIEF, and JUSTIFICATION.

Using **Cartesian knowledge** for EVIDENCE, we get SKEPTICAL CONSTRAINT:

SKEPTICAL CONSTRAINT 'S knows that P' is true at a world w and context c just in case for the c -determined constraint $\Psi^?$, every possibility consistent both with what S **knows** and with the correct answer to $\Psi^?$ (at the relevant world) is a P possibility.

8.2 Accounting for the data

Setting aside the data of §§1, 2, and 7 for the moment, one might wonder how various *ordinary* knowledge ascriptions are supposed to come out true on an account like SKEPTICAL

But given that, intuitively, I know that Jane might be work, shouldn't this sentence have problematic implicatures if IGNORANCE is true?

Not obviously. Following Dorr and Hawthorne (2013), we have reason to believe epistemic modals like 'might' take constrained readings too. The 'might' that appears in (J)'s consequent is constrained by the answer to the question 'Is Jane at home?'; and you *don't* know whether Jane might be at home in that sense of 'might'. One way of seeing that 'might' can take this reading is through closely related speeches like:

(J*) I don't know whether Jane might be at work; if she's at home then she definitely isn't, but if she's not at home then she might be.

The felicity of (J*) shows that there are resolutions of 'might' according to which I do not know whether Jane might be at work. Thus, (J)'s appropriateness is no problem for IGNORANCE.

³²We can accept that facts about appearances are not always "open to view" without impugning the legitimacy of the **knowledge** relation. See, e.g., Williamson (2000, ch. 4) and Schwitzgebel (2008) for arguments that we sometimes fail to know (or even just get wrong) facts about how things appear to us.

CONSTRAINT. Consider, for instance:

- (26) I know that I have hands.
(29) I know that Trump is the president of the United States.

Given the nature of **knowledge**, the propositions that one has hands and that Trump is the president are true at only relatively small subregions of **epistemic possibility**. So absent a constraint that allows us to ignore the rest of **epistemic space**, SKEPTICAL CONSTRAINT will have a difficult time vindicating our judgments about a number of platitudes about rather ordinary ascriptions.

Intuitively, then, what we need is a way of getting from **knowledge** to *knowledge*. And to do that we need a proposition that, if true, could take the set of possibilities consistent with what S **knows** and eliminate all and only those that are inconsistent with what she *knows*. If we had such a proposition, then we could form a question out of it and posit that as the “default” constraint when interpreting ordinary knowledge ascriptions.

To find the $\Psi^?$ that gets us from **knowledge** to *knowledge*, we will make use of another black box of sorts. Let N be the proposition that things are *normal*, where the notion of ‘normality’ is understood such that the proposition that things are normal is roughly equivalent to the proposition that things are more or less as they seem, or—another translation—to the proposition that one is in the good case. N is a theoretical posit, and a nice way of thinking about it is as whatever proposition the skeptic would have to be assured of before she would take herself to know the sorts of things a friend of *ordinary knowledge* thinks we know.³³

Given N, a natural thought is that *knowledge* is the relation you get when you constraint **knowledge** by $N^?$. Using our notation from earlier, *knowledge* is $\text{knowledge}_{N^?}$. Assuming the answer to $N^?$ is N, SKEPTICAL CONSTRAINT predicts that ‘S knows that she has hands’ is true in the *ordinary* sense of ‘knows’ just in case S $\text{knows}_{N^?}$ that she has hands. And this is indeed something S $\text{knows}_{N^?}$. Why? Well, although S doesn’t **know** that she has hands, S **knows** that it *appears* that she has hands. S also **knows** that if things are normal, then such appearances are veridical. S thus **knows** that if things are normal, then she has hands. This means that, for S, every **epistemic possibility** in which N is an **epistemic possibility** in which she has hands. So there are no **epistemic possibilities** consistent with the answer to $N^?$ (namely: N) in which she doesn’t have hands. Thus, S $\text{knows}_{N^?}$ that she has hands. More generally, whenever S **knows** a material conditional of the form ‘ $P \supset Q$ ’, then any constraint $\Psi^?$ whose true answer entails P will be such that S $\text{knows}_{\Psi^?}$ that Q.

So, *ordinary knowledge* is **Cartesian knowledge** constrained by the true answer to the question $N^?$. But what about the decidedly non-*ordinary*, non-**Cartesian** uses of ‘knows’ that feature in §§1-2’s (1)–(12)? Well, we know how to capture these uses in terms of

³³ It is not essential that N is a proposition rather than something like a *character*—i.e., a placeholder for whichever proposition the sentence ‘Things are normal’ expresses in context. Indeed, for reasons I lack the space to develop fully, I believe the most plausible interpretations of N are contextualist interpretations.

ordinary knowledge—this is what MODERATE CONSTRAINT does—and we just saw how to capture *ordinary knowledge* in terms of **Cartesian knowledge**. So to get SKEPTICAL CONSTRAINT to account for (1)–(12), we just need to assume that on their natural readings the contextually determined constraint is some combination of the question $N^?$ and whatever question MODERATE CONSTRAINT took the constraint to be in each instance. In the case of *History Exam*, for instance, the constraint is presumably some combination of the questions $N^?$ and ‘Does Peggy have good answers?’. Any constraint along these lines will be able to prune off all the **epistemic possibilities** in which the proposition that the answer to question 5 is b is false, thereby capturing (1)–(5). A similar strategy may be employed to translate MODERATE CONSTRAINT’s explanation of (6)–(12) into terms more friendly to SKEPTICAL CONSTRAINT.

Finally we may turn to the data that motivated us to move away from MODERATE to SKEPTICAL CONSTRAINT in the first place. With regards to §7.1’s skeptical puzzles, the story is straightforward. Let us focus on the (26)–(28) triad:

- (26) I know that I have hands.
- (27) If I know that I have hands, then I know that I am not a handless BIV.
- (28) I do not know that I am not a handless BIV.

Like MODERATE CONSTRAINT, SKEPTICAL CONSTRAINT gets CLOSURE for free (see footnote 29). So we just need to find plausible constraints to capture the intuitions about (26) and (28). Here’s the easy explanation: the constraint associated with the natural readings of (26) is $N^?$, the question of whether things are normal; while the constraint associated with the natural readings of (28) is either nothing—and so ‘knows’ denotes **knowledge**—or it is some other constraint sufficiently weaker than $N^?$ so as to leave behind **epistemic possibilities** in which one is a handless BIV.³⁴

The story for §7.2’s IGNORANCE implicatures is essentially the same, just now the question is how to interpret the ‘knows’ that appears in IGNORANCE rather than in sentences. But basically: the contexts in which (e.g.) (32) and (35) have unproblematic implicatures

- (32) If Oswald didn’t kill Kennedy, then someone else did.
- (35) If you are in the good case, then you know whether you have hands.

are those in which the ‘knows’ of IGNORANCE denotes a relation like **knowledge**. On many of their natural uses (32) and (35) indeed have the standard implicatures, it’s just that it’s

³⁴ Absent more detail about the context in which (28) is asserted it will be difficult to assess whether the constraint is empty or merely weaker than $N^?$. The point is especially vivid in the case of (31) (‘I do not know that Donald Trump has not had a fatal heart attack in the past few moments’). (31) is natural in a context in which perception of one’s environment suffices for knowledge of various local propositions—a context in which (28) is false—but in which non-local, contingent propositions are such that one can do no better than make educated guesses. Such a context is one in which the interpretation of ‘knows’ falls somewhere between **knowledge** and $\text{knowledge}_{N^?}$.

ignorance rather than *ignorance* that is implicated.

SKEPTICAL CONSTRAINT thus looks to be a strict improvement over MODERATE CONSTRAINT, at least for those who want to be in the business of treating our intuitions about skeptical knowledge ascriptions and ignorance implicatures as semantically respectable. And given the uniformity with which SKEPTICAL CONSTRAINT accounts for the data of §§1, 2, and 7, the presumptive view should be to treat it all on a par. As such, for the remainder of the paper I will suppose that SKEPTICAL CONSTRAINT is the superior implementation of SKELETON ACCOUNT.

9 The principles, revisited (again)

One straightforward consequence of adopting SKEPTICAL over MODERATE CONSTRAINT is that §6's "fallback" strategy for saving CREDENCE, BELIEF, and JUSTIFICATION from cases like *History Exam* begins to look a lot less plausible. There the strategy was to restrict the scope of the principles to *unconstrained* resolutions of 'knows', since on those resolutions the principles are valid (regardless of whether we go for MODERATE or SKEPTICAL CONSTRAINT).

But if SKEPTICAL CONSTRAINT is our preferred theory of 'knows', then the only unconstrained epistemic relation is **knowledge**. And although **knowledge** does satisfy each of CREDENCE, BELIEF, and JUSTIFICATION, the manner in which it does so provides at best a Pyrrhic victory for many of those who wish to defend them. Slogans like 'knowledge requires belief', 'knowledge requires justification', etc., are true only in virtue of the fact that **Cartesian knowledge**—i.e., the skeptical invariantist's knowledge—requires belief, justification, etc. But once one leaves the realm of the Cartesian, the principles cease to be robust generalizations. So if the question of the validity of CREDENCE, BELIEF, and JUSTIFICATION is the question of whether they are true in every context, then plausibly not one of the principles is valid. And if instead the question of their validity is the question of whether they are true on all *unconstrained* readings, then although each of the principles is valid, their validity isn't what we might have hoped for.

Is there some other theoretically interesting sense in which they might be valid? That is, is there some collection of privileged resolutions of 'knows'—where privilege is not a matter of the presence or absence of constraint—such that: the question of whether one of the principles is valid is tantamount to the question of whether the principle is true on all and only the privileged resolutions? I find the suggestion dubious. As far as I can tell, the division of labor among the resolutions of 'knows' countenanced by SKEPTICAL CONSTRAINT is too broad to lend much credibility to the thought that only certain special resolutions matter when assessing principles like CREDENCE or JUSTIFICATION. But if there were to be a special resolution, the obvious candidate would be knowledge_N ? (i.e., *knowledge*). So perhaps the principles are true in the way that matters because they are true of knowledge_N ?

(and **knowledge**), regardless of the fact that $\text{knowledge}_N?$ is just one of many constrained readings ‘knows’.

Now unfortunately a proper discussion of the merits of this view would take us too far afield, so instead I’ll have to settle for saying something suggestive about why I’m skeptical of its prospects. The plausibility of the envisioned response seems to presuppose that we are confident that our talk of *knowledge* picks out a *natural* epistemic relation, one whose nature we can reverse-engineer through attention to its various theoretical roles. But I think there are good reasons for skepticism that this is actually the case. For one, SKEPTICAL CONSTRAINT and its variety of non-*knowledge* epistemic relations is capable of doing much of the work *knowledge* was supposed to do. And for the stuff this package of relations doesn’t obviously handle—the trouble cases being the norms of rationality, evidential probabilities, and so on—it is entirely unclear whether *knowledge* does any better.³⁵ And against thinking that we’ve singled out a unique epistemic relation in our talk of *knowledge*, there are the familiar facts of the immense malleability in our judgments about what is known in cases in which we are supposed to be tracking *knowledge*. Here I have in mind (e.g.) Harman (1968)’s lottery paradox, Gendler and Hawthorne (2005)’s observations about fake barn cases, and Dorr et al. (2014)’s puzzle about our *knowledge* of coin flips.³⁶

10 Some over-generation worries

SKEPTICAL CONSTRAINT is designed to be flexible enough to accommodate the full range of ordinary uses of ‘knows’. Taking the judgments discussed by this paper at face value, it is truly a wide range: sometimes not even propositions about one’s immediate environment count as known, while other times propositions about opaque matters on which one has wildly guessed do. SKEPTICAL CONSTRAINT achieves its flexibility by having the thing that plays the EVIDENCE ROLE—**Cartesian knowledge**—include only a very limited set of propositions, and then allowing context to do all the work of eliminating whichever false propositions need to be eliminated to get an agent to know what we want them to know.

Of course, with gains in flexibility come losses in predictive power. Given the unruly nature of the data surveyed in this paper, we should probably learn to live with the fact that no empirically adequate theory of ‘knows’ is going to be all that predictive. Still, it’s one thing for a theory to predict that ‘S knows that P’ has both true and false readings (depending on how the context is resolved), yet not be able to tell us when we should expect to hear those readings; it’s another thing to predict that it has both true and false

³⁵ See, e.g., Bacon (2014) and Anderson and Hawthorne (forthcoming) for discussion on these points.

³⁶ And this is not even to begin to start on the massive literatures on pragmatic encroachment. For a sampling, see, e.g., Buckwalter and Schaffer (2015); Fantl and McGrath (2002, 2009); Hawthorne (2004, §4.2); May et al. (2010); McGrath (2007); Nagel (2008); Roeber (2018); Schaffer (2005a); Stanley (2005); Weatherson (2012).

readings when we can't seem to hear one of those readings at all. Theories with the first feature have the drawback of only being able to explain so much; theories with the second are empirically inadequate.

This section will present some worries of the second sort that naturally arise for views as bare-boned as SKEPTICAL CONSTRAINT.³⁷ The following section will then amend the view to avoid the problematic predictions.

According to SKEPTICAL CONSTRAINT, 'S knows that Q' is true at world w and context c just in case for the c -determined $\Psi^?$, S **knows** _{$\Psi^?$} that Q. Absent restrictions on which constraints are possible, the theory predicts that for *any* Q such that Q is true at w , there is a resolution of 'knows' in which S knows that Q. For example: the view predicts that for any P such that S **knows** that P entails Q, SKEPTICAL CONSTRAINT guarantees that S **knows** _{$P^?$} that Q.³⁸

Is this much flexibility really called for? Yes—or at least something close to it. Recall from the discussion of §5 that for most any **knows** _{$P^?$} relation, we can get a case where it is plausible that 'knows' denotes it. Again, assuming I **know** that P entails Q, any of the following will do the trick:

- (36) If P, then I know whether Q.
- (37) Either $\neg P$ or I know whether Q.
- (38) I might know whether Q; depends on whether P.

So that's something in favor of a view with as much bare-boned flexibility as SKEPTICAL CONSTRAINT.

On the other hand though, there are reasons to think that there are restrictions on the range of possible readings that are not reflected in SKEPTICAL CONSTRAINT. Here I have in mind three kinds of cases.

10.1 First worry: asymmetries in conditional ascriptions

As was already alluded to in §5, the fact that the subject term in each of (36)–(38) is a first-person pronoun seems to matter. Generally it is pretty easy to hear true readings of these first-person constructions. But instances of (36), for instance, often sound much worse when the subject term denotes someone who is not the speaker and who has no special evidence or beliefs about the antecedent and consequent. This can be seen in the following minimal pairs:

³⁷ Note that *any* theory of 'knows' that respects the basic structure of SKELETAL ACCOUNT (without substantial amendments) will face these problems, including MODERATE CONSTRAINT.

³⁸ Recalling §5, we know that there is some sort of restriction ruling out constructions of the form 'If P, then I know whether P'.

- (39) a. ✓ If Trump had a fatal heart attack only moments ago, then I know whether Pence is currently the president.
- b. ? If Trump had a fatal heart attack only moments ago, then Saul Kripke knows whether Pence is currently the president.
- (40) a. ✓ If Federer will win Wimbledon this summer, then I know how many times he'll have won Wimbledon by the end of the year.
- b. ? If Federer will win Wimbledon this summer, then my friend Kyle knows how many times he'll have won Wimbledon by the end of the year.

Let us focus on (39). The natural inference with (39b) is that Kripke is with Trump right now, and thus would know whether Trump has had a heart attack. Since we have no reason to think Kripke is or ever would be with Trump, (39b) seems false. By contrast, no such inference is required with (39a). It seems true even if we know that I wouldn't have any idea whether Trump has had a fatal heart attack in the past few moments. In other words: *my* *knowledge* of the fact that either Trump is alive or Pence is president seems to suffice for the existence of true readings of (39a). But for some reason *Kripke's* *knowledge* of the same doesn't seem to suffice for true readings of (39b). The problem is that nothing about SKELETON ACCOUNT (on either its MODERATE or SKEPTICAL implementations) is in a position to explain this fact.³⁹

10.2 Second worry: asymmetries in third-person ascriptions

Recall from *History Exam* that Pete's (bad) answers say the answer to question 5 is c, and that Pete has no idea what Peggy's answers say the answer to question 5 is. Recall also that Roger the TA eventually finds out that Peggy, not Pete, has the good answers. Now compare the following two utterances from Roger:

- (5*) ✓ Peggy knows whether the answer to question 5 is b.
- (5+) ? Pete knows whether the answer to question 5 isn't c.

As we know from §1, there are plenty of natural readings on which (5*) is true. But there seem to be no natural readings on which (5+) is.

We saw in §4.3 why according to the MODERATE implementation of SKELETON ACCOUNT,⁴⁰ (5*) is true when the contextually determined constraint is 'Does Peggy have the good answers?': Peggy has the good answers, she *knows* that if she has the good answers, then the answer to question 5 is b, and so every one of Peggy's *epistemic possibilities* consistent

³⁹ I'll note in passing that the contrasts seen in (39) and (40) are yet further evidence against the wide-scoping response to the data of §3.1. If the explanation of why (39a) is fine is that 'knows' is taking wide-scope over the entire conditional, then there is no reason (39b) shouldn't be fine too.

⁴⁰ (5*) and (5+) are analyzed in terms of MODERATE CONSTRAINT because it makes the problem easier to grasp. The same problem goes through on SKEPTICAL CONSTRAINT.

with the true answer to the question ‘Does Peggy have the good answers?’ is a *possibility* in which the answer to question 5 is b. Hence (5*).

The issue is that on any implementation of SKELETON ACCOUNT, (5+) should be true here too. And that is because Pete *knows* that if Peggy has the good answers, then the answer to question 5 *isn't* c. The answer to the constraint question is that Peggy has the good answers. Thus, all of Pete’s *epistemic possibilities* consistent with the true answer to the constraint question are *possibilities* in which the answer isn’t c. This gives us (5+). More generally, given the relevant symmetries in Peggy and Pete’s *evidence*—at least as concerns the exam—it is not clear there is *any* constraint that makes (5*) true but (5+) false. But again this prediction is not borne out by the data.

10.3 Third worry: true/false exams

A similar worry arises if we reimagine *History Exam* as involving a true/false rather than a multiple choice exam. As before, Peggy and Pete *know* that their answer sheets disagree on the answer to every single question. So, for example, if question 5 is ‘Is it true or false that the Berlin Wall fell in 1989?’, and Peggy’s answer sheet says ‘True’, she’ll *know* that Pete’s says ‘False’. What this means is that now both students are in a position to *know* what the other student’s answer sheet says (whereas on the multiple choice version they could only guess with $\frac{1}{3}$ confidence).

My intuition is that with the exception of (1)—more on which in the next section—the *History Exam* ascriptions sound noticeably degraded on this version of the case:

- (1) ✓ If my answer sheet is good, then I know what the answer to question 5 is.
- (2) ? One of us knows what the answer to question 5 is.
- (3) ? I might know what the answer to question 5 is; it depends on whether I got the good answers.
- (4) ? Whoever has the answer sheet that says ‘5: True’ knows the answer to every question.
- (5) ? Peggy knows what the answer to question 5 is.

Why the difference between the multiple choice and true/false versions of *History Exam*? Intuitively it is because on the true/false version of the case Peggy and Pete have, for all intents and purposes, shared their answers. They each *know* what both of their answer sheets say for every question. So, intuitively, for any given question, either both should know the answer or neither should. The possibility that one could know without the other knowing is hard to make sense of given the intuitive symmetry of their epistemic states.

But none of this matters as far as SKELETON ACCOUNT is concerned. As readers can see for themselves, (1)–(5) still come out true on the true/false version of the case when the

constraint question is ‘Does Peggy have the good answers?’. So if the proponent of SKELETON ACCOUNT (on either its MODERATE or SKEPTICAL implementations) wants to explain why (1)–(5) seem worse for the true/false exam, she’s going to have to rely on a theory of the meta-semantics of constraint to do so. That is to say, she’s going to have to argue that for some reason or other, shifting from a multiple choice to a true/false version of *History Exam* makes the relevant constrained readings harder to hear. The prospects of this line seem dubious at best.

We have three over-generation challenges that look like they should apply to just about any implementation of SKELETON ACCOUNT (whether MODERATE or SKEPTICAL). First, the theory doesn’t tell us why there is an intuitive difference between first- and third-person conditional knowledge ascriptions. Second, the theory doesn’t tell us why Peggy can know the answer to question 5 is b while Pete seemingly can’t know the answer isn’t c. And third, the theory doesn’t tell us why the ascriptions start to sound worse once we move from a four-question multiple choice exam to a true/false exam. The aim of the next section is to amend SKELETON ACCOUNT in a way that answers these challenges.

11 Knowing and thinking

I believe these three challenges admit of a uniform answer, though the idea behind it is complicated and deserves more space than can be given here.⁴¹ So what I say will have something of the form of an outline of a response, rather than a response proper. But I think there should be enough to show how the proponent of SKELETON ACCOUNT—so-called for a reason!—has the resources to substantively reign in some of the view’s more problematic predictions.

11.1 ‘Knows’ and ‘thinks’

Most of the paper has focused on our intuitive judgments about ‘knows’ in cases like *History Exam*. But something interesting happens with analogous constructions involving ‘thinks’.⁴² For starters: notice that in the context of the original multiple choice version of *History Exam*, it is natural to claim that each of the students *thinks* the answer to question 5 is the one listed on his or her answer sheet:

(41) a. ✓ Peggy thinks the answer to question 5 is b.

⁴¹ See Holguín (2018a) for a full-length discussion of the view defended in this section.

⁴² This section focuses on ‘thinks’ rather than ‘believes’. This is because first, the ‘thinks’ sentences it discusses are in general somewhat better than the analogous ‘believes’ sentences. And second, because our judgments about ‘thinks’—an ordinary bit of language—are less likely to be tainted by philosophical preconceptions than our judgments about ‘believes’—a quasi-technical bit of language. We will return to the connection between ‘thinks’ and ‘believes’ at the end of the section.

- b. ? Peggy thinks the answer to question 5 isn't b.
- (42) a. ✓ Pete thinks the answer to question 5 is c.
- b. ? Pete thinks the answer to question 5 isn't c.

One might wonder whether the difference in our judgments about (41a) and (42b) might have something to do with the difference in our judgments about §10.2's (5*) and (5+):

- (5*) ✓ Peggy knows whether the answer to question 5 is b.
- (5+) ? Pete knows whether the answer to question 5 isn't c.

Supposing it does, we have reason to believe that THINKING is true:

THINKING If S knows that P, then S thinks that P.

And by incorporating THINKING into SKELETON ACCOUNT, we would have the resources to block at least the challenge from (5*) and (5+). In particular, we could say that because Pete *thinks* the answer to question 5 is c (when in fact it's b), no matter the *c*-determined constraint he cannot know whether the answer to question 5 isn't c.

Let's continue to run with the THINKING thought for the other two challenges. Consider the 'thinks' analogs of §10.1's (39) and (40):

- (43) a. ✓ If Trump had a fatal heart attack a few moments ago, then I think Pence is currently the president.
- b. ? If Trump had a fatal heart attack a few moments ago, then Saul Kripke thinks Pence is currently the president.
- (44) a. ✓ If Federer will win Wimbledon this summer, then I think he'll have won Wimbledon nine times by the end of the year.
- b. ? If Federer will win Wimbledon this summer, then my friend Kyle thinks he'll have won Wimbledon nine times by the end of the year.

The judgments concerning 'knows' and 'thinks' seem to continue to pair with each other in these cases, suggesting that the contrast between first- and third-person conditional ascriptions can be accounted for by THINKING as well.

Finally, consider the 'thinks' analogs of *History Exam's* (1)–(5) on the original *multiple choice* version of the case:

[In Peggy's mouth:]

- (45) ✓ If my answer sheet is good, then I think the answer to question 5 is b.
- (46) ✓ One of us thinks the answer to question 5 is b.

(47) (?) I might think the answer to question 5 is b.

[In Roger's mouth:]

(48) ✓ Whoever has the answer sheet that says '5: b' thinks the answer to question 5 is b.

(49) ✓ Peggy thinks the answer to question 5 is b.

With the exception of (47), each of these 'thinks' ascriptions sounds fine. So too do the original 'knows' ascriptions (1)–(5).⁴³

By contrast, with the exception of the conditional 'thinks' ascription (50), none of the 'thinks' analogs of (1)–(5) sounds fine on §10.3's *true/false* version of *History Exam*:

[In Peggy's mouth:]

(50) ✓ If my answer sheet is good, then I think the answer to question 5 is True.

(51) ? One of us thinks the answer to question 5 is True.

(52) ? I might think the answer to question 5 is True.

[In Roger's mouth:]

(53) ? Whoever has the answer sheet that says '5: True' thinks the answer to question 5 is True.

(54) ? Peggy thinks the answer to question 5 is True.

And as we know from §10.3, only the conditional 'knows' ascription is acceptable on the *true/false* version of *History Exam*. We thus see the following pattern: on the multiple choice version of *History Exam*, both the 'knows' ascriptions and 'thinks' ascriptions sound fine; but on the *true/false* version, both sound problematic (with the conditional ascriptions being exceptions). Is there anything that can be said about why this pattern should arise?

Well, if THINKING and some descendant of SKELETON ACCOUNT are both true, then it is entirely unsurprising why the 'knows' and 'thinks' ascriptions should stand and fall together. The real question, then, is why our judgments about the 'thinks' ascriptions differ between the multiple choice and *true/false* versions of *History Exam*.

⁴³ The badness of (47) is plausibly due to the fact that Peggy *knows* that she thinks the answer to question 5 is b, and thus that it carries problematic implicatures (rather than that it is false). Observe that its 'certain' analog is fine:

(47*) ✓ It is certain that I think the answer to question 5 is b.

Here is a stab at an answer. On the true/false version of the case, there are only two possible answers to each question: True or False. Thus, for any given question Peggy *knows* what the answer is conditional on her answer sheet being good *and* conditional on her answer sheet being bad. Moreover, Peggy *knows* that she and Pete were equally likely to have received the good answers. So for any given question, Peggy *knows* it is no more likely that she'll get the right answer by answering in accordance with her sheet than it is that she will get it right by answering against it. Peggy thus has no reason to guess with her answer sheet rather than against it, and by extension has no basis on which to be opinionated about the answer to any given question. This explains why 'Peggy thinks the answer to question 5 is b' is false on the true/false version of the case. By contrast, on the multiple choice version of the case, Peggy *does* have reason to guess in accordance with her answer sheet. And that is because she *knows* that if her answer sheet is good, then the answer to any given question is the one it recommends, but if her answer sheet is bad, then the answer to any given question could be one of *three* possibilities. So Peggy *knows* she does better by guessing in accordance with her answer sheet than she does guessing against it. She thus has good reason to be opinionated in line with her answer sheet on each of the questions. As a result, 'Peggy thinks the answer to question 5 is b' is true on the multiple version of the case. And this, in turn, explains the intuitive contrast between (45)–(49) and (50)–(54).

The judgments here are subtle, but the pattern between the 'knows' ascriptions (embedded or unembedded) and the 'thinks' ascriptions is striking enough that THINKING ought to be taken seriously. There is thus a strong *prima facie* case for upgrading from SKELETON ACCOUNT to:

SILHOUETTE ACCOUNT 'S knows that P' is true at a world w and context c just in case: (i) 'S thinks that P' is true in c ; and (ii) for the c -determined constraint $\Psi^?$, every possibility consistent both with S's EVIDENCE and with the correct answer to $\Psi^?$ (at the relevant world) is a P possibility.

Given the judgments about 'knows' and 'thinks' outlined above, SILHOUETTE ACCOUNT is a strict improvement over SKELETON ACCOUNT. It seems to avoid the over-generation challenges described in the previous section without compromising the theory's empirical reach.

11.2 'Thinks' and 'believes'

Before closing I want to make two quick remarks about THINKING in the present context. First, even supposing we take all the above data at face value, nothing I have said so far gives us any indication of what the proper semantics for 'thinks' looks like. We know a few things: that 'S thinks that P' can be true even when both (i) 'S has credence greater than .5 (or even .001) that P' is false and (ii) 'S displays the normal behaviors associated with a

belief in P' is false; that 'S thinks that P' can be true on its natural readings when embedded in a conditional but false on its natural readings when unembedded; and that having equally good evidence in favor of P and \neg P seems to make it hard to hear the true readings of 'S thinks that P' (if there even are any). But a list of generalizations is far from a theory of the concept.

Second, supposing we like THINKING, one might wonder whether we have reason to revisit §1's arguments against BELIEF:

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

After all, it is difficult to deny the following principle:

THINKING IS BELIEVING If S thinks that P, then S believes that P.

As observed elsewhere in the literature, conjunctions of the form

(55) ? S thinks that P, but it's not that S believes that P.

are abominable.⁴⁴ But supposing THINKING and THINKING IS BELIEVING are both true, it follows that BELIEF is true as well. What then to make of §1's arguments against BELIEF?

Well, there were just two arguments there: first, that *History Exam* shows us that 'S knows that P' is sometimes compatible with 'S has arbitrarily low credence in P'; and second, that *History Exam* shows us 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 we can dispense of this background assumption and instead take the lesson to be that if BELIEF is true, then cases like *History Exam* give us reason to believe that 'believes' is capable of denoting some exotic doxastic relations. Supposing that one is willing to embrace this conclusion, then one need not think that any of the data considered in this paper gives us serious reason to doubt the validity of BELIEF. Instead one can think that the strange knowledge of §§1–2 shows us that there is also such a thing as strange belief.

12 Conclusion

I have argued that the ways in which we ordinarily talk and reason about knowledge are best explained by the mechanisms of constraint. Supposing we wish to account for rather than ignore these facts, we should think that *some* descendant of SKELETON ACCOUNT—and in particular one of its SILHOUETTED descendants—has probably got to be the correct theory of knowledge. But given the nature of cases like *History Exam*, *Memory Experiment*, and *Battery*, we also know that whichever account emerges as the best contender, widely accepted

⁴⁴ See Hawthorne et al. (2016).

principles about knowing are going to have to be reconsidered. CREDENCE and JUSTIFICATION look to be more or less unrecoverable—at least if the relevant notions are intended to play their standard theoretical roles. And while BELIEF may be recoverable after all, we'll have to get used to thinking of belief as being compatible with arbitrarily low credence and the like.

But what about knowledge itself? What picture does a **Cartesian knowledge**-based implementation of SILHOUETTE ACCOUNT leave us with? Well, stated roughly, it is the view that knowledge is whatever is both believed and entailed by the conjunction of one's **evidence** and certain propositions selected by context—namely those that are the true answers to a collection of contextually determined questions. Since the truth is always compatible with one's **evidence**, it follows that for any true P , if one thinks that P then there's at least a sense in which one knows that P . Thus, in any given context c , the difference between one who knows that P and one who doesn't is that for the c -determined $\Psi^?$, the one who knows **knows** that the true answer to $\Psi^?$ entails P , while the one who fails to know doesn't **know** that the true answer to $\Psi^?$ entails P . Consequently, becoming an expert on a topic is just a matter of gaining **knowledge** of propositions of the form $\Psi \supset P$ (for Ψ that happen to be true).

As a result, it's not knowledge first, but **Cartesian knowledge** first—at least as concerns the analysis of knowledge. It is also a picture of knowledge that is inhospitable to knowledge first theories of the distinctively normative questions we find in epistemological inquiry: questions about what the difference is between rational and irrational beliefs, whether and how we are justified in making inductive inferences, how we ought to respond to cases of peer disagreement, and so on. The basic reason why is that any plausible descendant of SILHOUETTE ACCOUNT will posit a plurality of knowledge relations—indeed, one for every possible constraint question. Since none of these relations is obviously more joint-carving than the others (with perhaps the exception of whichever one we associate with the EVIDENCE role), one's pluralism about knowledge relations will force equally pluralistic answers to these other normative questions. But it is not clear whether pluralistic answers to these sorts of questions are wholly satisfactory. Take the question of whether we are justified in making inductive inferences. If one thinks this question goes deeper than the question of which $\Psi^?$ s combine with our **knowledge** to entail the reliability of induction—that it's important that we figure out whether these inferences are *justified*, rather than whether they are justified₁ or justified₂—then one should also think that the path to progress will not be found by thinking about how our knowledge of induction works.⁴⁵

But what the view sacrifices in knowledge first epistemology it gains in a knowledge first philosophy of language. A contextualist theory of 'knows' as flexible as SILHOUETTE

⁴⁵ I am not ruling out the possibility that going pluralist about justification relations—presumably to match SILHOUETTE ACCOUNT's pluralism about knowledge relations—is the best we can do to answer these sorts of questions. But it does seem preferable to avoid that path if possible.

ACCOUNT has the resources to account for all sorts of data thought to be recalcitrant for knowledge first theses concerning language and communication: knowledge-theoretic accounts of assertion and ignorance implicatures, the semantics of epistemic vocabulary like ‘if’, ‘might’, ‘certainly’, and so on. Cases that have been taken to put pressure on the links between these concepts and the concept of knowledge can now more readily be explained away as invoking non-uniform resolutions of the context-sensitivity of ‘knows’. All that is required is a commitment to the view that the norms of assertion, the pragmatics of indicative conditionals (and disjunctions), and the semantics of epistemic vocabulary are all roughly as context-sensitive as ‘knows’ is.⁴⁶ Whether this is the kind of commitment we should be want to take on is a question for future work. The important point is just the abstract one: if we adopt a theory of knowledge as flexible as the one defended in this paper, we will have a more difficult time using the concept of knowledge to answer questions about how we ought to think, but an easier time using it to answer questions about how we in fact think. This seems to me a nice feature of the view.

⁴⁶ See [Worsnip \(2017\)](#) and [Holguín \(2018b\)](#) for discussion of some of the issues involved with pairing contextualism about ‘knows’ with various kinds of knowledge norms. And see also [Blumberg and Holguín \(2018\)](#) for a discussion of how the mechanisms of constraint can help explain some other puzzling data involving emotive factive verbs like ‘surprise’ and ‘regret’.

References

- Anderson, C. and Hawthorne, J. (Forthcoming). Knowledge, practical adequacy, and stakes. In Gendler, T. and Hawthorne, J., editors, *Oxford Studies in Epistemology*. Oxford University Press.
- Bacon, A. (2014). Giving your knowledge half a chance. *Philosophical Studies*, (2):1–25.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118(470):241–294.
- Blumberg, K. and Holguín, B. (2018). Embedded attitudes. *Journal of Semantics*, pages 1–38. Forthcoming.
- Buckwalter, W. and Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs*, 49(2):201–234.
- Cohen, S. (1986). Knowledge and context. *Journal of Philosophy*, 83(10):574–583.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2(n/a):91–123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13(s13):57–89.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4):913–929.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1):1–52.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111(2):167–203.
- Dorr, C., Goodman, J., and Hawthorne, J. (2014). Knowing against the odds. *Philosophical Studies*, 170(2):277–287.
- Dorr, C. and Hawthorne, J. (2013). Embedding epistemic modals. *Mind*, 122(488):867–914.
- Fantl, J. and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1):67–94.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford University Press.
- Farkas, K. (2016a). Know-wh does not reduce to know-that. *American Philosophical Quarterly*, 53(2):109–122.
- Farkas, K. (2016b). Practical know-wh. *Noûs*, 50(2).
- Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition and Logical Form*. New York: Academic Press.
- Gendler, T. S. and Hawthorne, J. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124(3):331–352.
- Gillies, A. (2010). Iffiness. *Semantics and Pragmatics*, 3(4):1–42.

- Greenough, P. and Kindermann, D. (2017). The semantic error problem for epistemic contextualism. In Ichikawa, J., editor, *Routledge Handbook of Epistemic Contextualism*, pages 305–320. Routledge.
- Harman, G. (1968). Knowledge, inference, and explanation. *American Philosophical Quarterly*, 5(3):164–173.
- Hawthorne, J. (2000). Implicit belief and a priori knowledge. *Southern Journal of Philosophy*, 38(S1):191–210.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In Steup, M. and Sosa, E., editors, *Contemporary Debates in Epistemology*, pages 26–43. Blackwell.
- Hawthorne, J. (2007). Eavesdroppers and epistemic modals. *Philosophical Issues*, 17(1):92–101.
- Hawthorne, J., Rothschild, D., and Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5):1393–1404.
- Holguín, B. (2018a). Believing is guessing. Unpublished manuscript.
- Holguín, B. (2018b). Knowledge in the face of conspiracy conditionals. Unpublished manuscript.
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155(3):383–398.
- Klein, P. (2015). Skepticism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2015 edition.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567.
- May, J., Sinnott-Armstrong, W., Hull, J. G., and Zimmerman, A. (2010). Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1(2):265–273.
- McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3):558–589.
- Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 5(8):1–81.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2):279–294.
- Neta, R. (2003). Contextualism and the problem of the external world. *Philosophy and Phenomenological Research*, 66(1):1–31.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69.
- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs*, (2):171–195.

- Rothschild, D. and Spectre, L. (2016). A puzzle about knowing conditionals. *Noûs*.
- Rothschild, D. and Spectre, L. (2017). At the threshold of knowledge. *Philosophical Studies*, pages 1–12.
- Schaffer, J. (2005a). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies*, 127(1):87–107.
- Schaffer, J. (2005b). What shifts? : Thresholds, standards, or alternatives? In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.
- Schaffer, J. and Szabo, Z. G. (2013). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*, (2):1–53.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96(1):317–333.
- Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review*, 117(2):245–273.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Weatherson, B. (2012). Knowledge, bets, and interests. In Brown, J. and Gerken, M., editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219):213–235.
- Worsnip, A. (2017). Contextualism and knowledge norms. In Ichikawa, J., editor, *The Routledge Handbook of Epistemic Contextualism*. Routledge.