Weak Knowledge

Sam Carter

Abstract

This paper is about the connection between your credences and what you know. After setting the scene (§§1-2), I argue that this connection is less demanding than is often thought: knowing something is compatible with assigning it any credence other than zero. Building on recent work in Holguín (2022) and Dorst & Mandelkern (2023), I propose (§3) you can know on the basis of a good guess.

Having outlined the view, I turn to defending it. I'll argue that it provides a better fit with out intuitions than competitors (§3.1) and—when combined with an appropriate package of norms on assertion—correctly predicts our use of knowledge ascriptions (§3.2). In the remainder of the paper (§§4-6), I'll draw out some additional consequences of the view and respond to objections.

1 Introduction

What does it take to know? According to a common consensus, knowledge requires a doxastic commitment to what is known; in order to know, you must—among other things—take a position on the issue you know about. Doxastic commitment, in the broad sense intended, comes in different forms. You can doxastically commit to a proposition by believing it, predicting it, being confident of it or being sure of it (among other attitudes).

My main focus in this paper is on what strength of commitment is required for knowledge. Specifically, I'm going to argue that the requirement is significantly less demanding than is typically assumed. All that is necessary for an agent to satisfy it is that they make a good guess. In this respect, the present proposal bears a family resemblance to recent work arguing that the strength of commitment required for belief is (much) weaker than has typically been thought (Hawthorne *et al.* (2016); Dorst (2019); Rothschild (2020); Holguín (2022); Dorst & Mandelkern (2023)).

The proposal in this paper has a number of distinctive features: it implies that you can know something despite rationally assigning it an arbitrarily low credence; that you can know something even when your best guess is that you do not know it; and that knowledge can fail to be luminous for reasons independent of margins for error.

Here's the plan. The central thesis is introduced and motivated in §3. I discuss some of its advantages, drawing on evidence of various different kinds. §4 raises and addresses a range of objections. §5 introduces a simple model of the theory and looks at the predictions it makes about higher-order knowledge. §6 concludes. To start with, however, we need to consider what it is for an agent to make a good guess. §2 offers a brief summary of recent work on this issue.

2 Guessing

Imagine an election with four candidates: A, B, C and D. According to polling data, each candidate has the following chance of winning:

A	B	C	D
40%	25%	20%	15%

Holguín (2022) and Dorst & Mandelkern (2023) observe that cases like this give rise to an interesting pattern of judgments. Whereas (1.a-c) would be rational guesses about the question *Who will win the election?*, (2.a-c) would not.[1]

(1) a. A will win.
 b. A or B will win.
 c. D won't win.

(2) a. B will win.
 b. A or C will win.
 c. A won't win.

There are a few noteworthy features of this pattern. First, it implies a guess may be rational despite being (i) less likely than an irrational guess and/or (ii) less likely than its negation. For example, (1.a) is a good guess, but (2.c) is a bad guess. Second, it implies a guess may be irrational despite being entailed by a rational guess. For example, (2.b) is a bad guess, but (1.a) is a good guess. Finally, it implies that there may be more than one way of guessing rationally. For example, each of (1.a-c) would be good guesses.[2] All of these features are succinctly captured on the account(s) of good guessing proposed by Holguín and Dorst & Mandelkern.

Guesses are directed at questions. An answer can be a good guess about one question while being a bad guess about another. Accordingly, we want an account of good guessing which is question-relative. We can think of a question as a partition on a set of worlds (Hamblin (1958, 1973); Groenendijk & Stokhof (1984)). Each cell of the partition can be thought of as corresponding to a complete answer. Complete answers are propositions which settle everything about the way the world is which is relevant to the question. The closure of the set of cells under union corresponds to the set of partial answers (which includes the complete answers as a limiting case). Partial answers are propositions which do not settle anything about the way the world is which is irrelevant to the question. Where Pr is a classical probability measure, we start by defining the notion of a cogent answer to a question.

[1] Here, and throughout, I use italicization to denote the content (e.g., question/proposition) expressed by the sentence italicized.

[2] Which answer an agent in fact guesses is often the outcome of a volitional act. As Holguín emphasises, it does not follow that it is always under an agent's control what guess they make. Perceiving or remembering the answer to a question may entail a commitment to that answer being correct which suffices to count as having a guess.

Cogency For any question, Q, p is a cogent answer to Q (relative to Pr) iff there is some non-empty $X \subseteq Q$ such that:
(i) $p = \bigcup X$; and
(ii) For all $q \in X$ and $r \in Q - X$: $Pr(q) > Pr(r)$.

Consider some partial answer to a question. That answer is cogent iff every complete answer which is compatible with it is strictly more probable than any complete answer which is incompatible with it. Here's another way of getting at the same idea: say that two partial answers are equally strong iff they are compatible with the same number of complete answers to a question. Then, for any question with only finitely many complete answers, an answer is cogent iff it is strictly more probable than every equally strong distinct answer. So, for example, each of (2.a-c) fails to be cogent because the corresponding answer from (1.a-c) is equally strong and more probable.

Holguín proposes to understand what it is for a guess to be rational (i.e., what makes a guess good) in terms of cogency.[3] Assume that an agent's evidence fixes whether the credence they assign to the partial answers to a question is rational. In this case, let's say that their credences are rational over the question.[4] Then, where Cr_S is a probability measure representing S's credences, we have:

Guessing p is a good guess about Q for S iff:
(i) p is a cogent answer to Q (relative to Cr_S); and
(ii) Cr_S is rational over Q.

Good guesses are cogent answers (relative to a rational way of assigning credence over the relevant question). For example, the cogent answers to the question *Who will win the election?* are: (i) A; (ii) A or B; (iii) A, B, or C; and (iv) A, B, C, or D. It is easy to check that the account correctly predicts the pattern of judgments observed above. An answer may be cogent despite being less probable than it's negation and/or less probable than some non-cogent answer; not every consequence of a cogent answer will be itself cogent; and there may be more than one cogent answer to a question. However, in addition to making satisfying predictions about the conditions under which a guess is rational, this account can also help to illuminate some distinctive features of knowledge.

[3] Dorst & Mandelkern (2023) argue instead for the following:
Guessing* p is a good guess about Q for S iff guessing p about Q maximizes expected utility relative S's epistemic utility function and Cr_S.
Dorst and Mandelkern show that, as long as S's epistemic utility function is *truth-directed* and *question-based*, then p will be a good guess only if it is a cogent answer to Q. However, **Guessing*** is more restrictive than **Guessing**, since some cogent answers may fail to be good guesses for S, depending on the epistemic utility function S adopts.

[4] According to **Guessing**, agents who assign irrational credences to a question's answers cannot make a good guess about that question (though they may still make good guesses about distinct questions). How demanding this constraint is will depend in part on whether there is exactly one rational assignment of credences given an agent's evidence (White (2005); Feldman (2007); Kelly (2013); Schoenfield (2014), a.o.). Those who deny permissivism may wish to weaken condition (ii), to require, for example, that Cr_S approximates the rational credence assignment over Q or that Cr_S orders the probability of complete answers to Q in the same way as the rational credence function.

3 Knowing

My primary goal is to argue that, when combined with the account of guessing above, the following provides a satisfactory account of knowledge:

Knowing S knows p relative to Q iff:
(i) p is an answer to Q;
(ii) S makes a good guess about Q which entails p;
(iii) S's guess about Q is correct and is not gettierized;

To know an answer to a question is to make a good guess about it which entails that answer and for the guess you make to be true and ungettierized. Every answer to a question entails itself. So, if an agent knows any answer to a question, the strongest answer they know will be their guess about it. Or to put it another way: an agent cannot know any answer to a question unless they know the answer they guessed.[5]

Knowing rules out—by fiat—knowledge in cases where an agent's guess is correct due to gettierization. In §5, I propose an operationalization of this condition in terms of reliability: roughly, the idea is that way an agent assigns credences to answers must reliably lead them to make true guesses. To avoid gettierization, it must be modally robust that the good guesses (for them) are correct guesses. To the extent that we should not expect to be able to grasp the relevant notion of reliability independently from our grasp of knowledge, we should not expect **Knowing** to serve as a basis for a non-circular analysis (cf. Carrier (1977), Fodor (1998, §§3-4), Williamson (2000, §1)).

Even if it will not serve a basis for a non-circular analysis of knowledge, the proposal makes a number of distinctive predictions. However, before considering these predictions, we need to say something about its relationship to knowledge ascriptions in natural language.

Knowing treats knowledge as a three-place relation, which holds between an agent, a proposition and a question. In this respect, it is a contrastivist account (cf. Karjalainen & Morton (2003); Morton & Karjalainen (2008); Schaffer (2004, 2005, 2007a,b, 2009, 2012); Schaffer & Knobe (2012); Schaffer & Szabó (2014); Hughes (2013); Sawyer (2014)). Although constrastivism about knowledge has been argued (most extensively by Schaffer) to have both theoretical and empir-

[5]Some authors have argued that agents can know a weaker, true proposition via doxastic commitment to a stronger false proposition (Hilpinen (1988); Williamson (2011, 2013, 2014)). This is sometimes motivated by the idea that knowledge should be subject to so-called graceful degradation: agents committed to an incorrect answer which approximates the correct answer know (strictly) more about a question than agents committed to an incorrect answer which is not even approximately correct. While **Knowing** lacks graceful degradation for knowledge, it retains a form of graceful degradation for what an agent is in a position to know. That's because agents are always in a position to know the strongest, true cogent answer to a question.

Arguably, this is all we should want. Our inability to know on the basis of incorrect guesses can help to explain why guessing a weaker answer may sometimes be preferable to guessing a stronger answer. Although more informative, the latter carries a greater risk of leaving us knowing nothing.

ical advantages, some care is necessary in how the view is stated.

The English verb 'knows' (like its equivalents in other languages) takes a single clausal complement: either a declarative clause (in so-called 'knowledge-that' constructions) or interrogative clause (in so-called 'knowledge-wh' constructions), but not both. Accordingly, any proponent of a contrastivist account owes an explanation of how knowledge ascriptions in ordinary talk express the three-place relation they posit.[6]

I'll assume that contexts associate propositions with questions; in any conversation, the context fixes, for each proposition, a unique salient question which that proposition answers. Salience is, in part, a matter of what is said in a conversation. There are various mechanisms by which a question's salience might be raised. For example, overtly asking it or using an attitude-verb which embeds it can be expected to make a question salient.

Following standard terminology, I'll say Q is a part of Q' iff every complete answer to Q is (at least) a partial answer to Q'. I'll use 'ϕ' as a schematic variable over declarative clauses; 'π' as a schematic variable over interrogative clauses; and 'α' as a schematic variable over names.[7] I adopt the following semantic clauses:

(a) $\ulcorner\alpha$ knows $\phi\urcorner$ is true in c iff for any Q salient in c: if the proposition ϕ is a partial answer to Q, then α knows ϕ relative to Q (at the world of c).

(b) $\ulcorner\alpha$ knows $\pi\urcorner$ is true in c iff for any Q salient in c: if Q is a part of the question π, then there is some $p \in Q$ such that α knows p relative to Q (at the world of c).

Knowledge-that ascriptions are true iff the subject knows the proposition expressed by the declarative clause relative to every contextually salient question to which it is a partial answer. Knowledge-wh ascriptions are true iff the subject knows a complete answer to every contextually salient question which is part of the interrogative clause.

To avoid vacuous quantification, we'll assume that knowledge-that constructions presuppose there is a salient question which the proposition expressed by the declarative clause partially answers. Similarly, knowledge-wh constructions presuppose there is a salient question which is part of the question expressed by the interrogative clause. In the simplest case, the unique such question will be the question expressed by the interrogative clause itself (since, by assumption, this question is raised to salience by the assertion). In this case, $\ulcorner\alpha$ knows $\pi\urcorner$ will be true iff α knows a complete answer to π. In what follows, I will primarily focus on examples of this kind. However, in §4.4, I'll consider cases in which

[6] Throughout what follows, I use 'knows' in a non-standard, technical way to refer this three-place relation between agents, propositions and questions. In contrast, where I mention 'knows' (in quotation or in numbered examples), it only ever refers to the standard verb of English, which takes either an interrogative or declarative.

[7] It follows that ϕ is a proposition—namely, the proposition expressed by $\ulcorner\phi\urcorner$. Similarly, π is a question—namely, the question expressed by $\ulcorner\pi\urcorner$.

this assumption needs to be dropped.

With these clauses in hand, we can now turn to consider a key property of the view, on which it departs from most alternatives (whether contrastivist or non-contrastivist).

According to **Knowing**, knowing an answer to a question is compatible with assigning it an arbitrarily low positive credence. More carefully, there is no $k > 0$ such that, for every Q, an agent must assign p a credence greater than k in order to know p relative to Q. In this sense, knowledge is weak: the form of doxastic commitment required for knowledge imposes no general requirement on an agent's credence in the proposition committed to, other than that it be greater than 0. Contrast with this views which take knowledge to be strong. A view treats knowledge as strong iff the form of doxastic commitment it takes to be necessary for knowledge requires an agent's credence to exceed some non-trivial positive threshold. Typically, this threshold will be at least greater than .5. Views on which knowing requires being confident, being sure, or being certain all take knowledge to be strong.

In saying that **Knowing** treats knowledge as weak, it is important to distinguish how reliably evidence supports correct answers from how strongly it supports them. An agent's evidence could reliably lead them to assign a question's true complete answer higher credence than any complete false answer despite only requiring that the credence they assign to it is slightly higher. For instance, consider an electoral poll (like the one in §2). It is perfectly coherent to suppose both that: (i) the polling data could not easily have made it rationally permissible to take some non-winner to be the most likely candidate to win but (ii) the polling data only rationally requires you to treat the winner as slightly more likely to win than each of the non-winners. Accordingly, if reliability of this kind suffices for non-gettierization, we have no reason to expect that the non-gettierization condition will impose independent constraints on the credence agents assign to what they know. The fact that knowledge precludes gettierization is not grounds for thinking knowledge is not weak.

In §§3.1-2, I'll argue we should prefer a view on which knowledge is weak to views on which knowledge is strong. I'll do this by showing that there are a range of observations which would be hard to explain if knowledge were strong, but which are expected if knowledge is weak.

3.1 Weakness

Consider, again, our four-candidate election. Imagine that, prior to the election, individuals are asked to guess who will win. Based on the pre-election polling data, Lia guesses it will be A. Imagine, further, that A in fact wins. Then (assuming she is not gettierized) **Knowing** predicts (3) and (4) have true readings.[8]

[8] At least, it does given the account of good guessing in **Guessing**. From here on, I'll take **Guessing** for granted when considering the predictions of **Knowing**.

(3) Lia knew who would win the election.

(4) Lia knew that it would be A who won.

Both of these predictions, I suggest, are correct. Not only are we are willing to ascribe knowledge to agents on the basis of (mere) true good guesses, such ascriptions are commonplace.

This point is supported by the kinds of response which are appropriate on learning that someone made a correct guess. Upon hearing the outcome of the election, someone aware of Lia's pre-election guess could reasonably ask 'How did you know that would happen?' (to which a reasonable response would be: 'I looked at the polling data and it ranked A as the favorite'). Yet 'how'-questions presuppose the truth of embedded clause. Asking 'How did John make $1,000,000 on the stock market?' is marked if it is not common ground (and cannot be easily accommodated) that John did make $1,000,000 on the stock market. Similarly, unless it was common ground (or could be easily accommodated) that (3)-(4) are true in the circumstances, we would not expect this form of response to be appropriate.

A range of further evidence supports the conclusion that knowledge is weak. First, in situations in which it can be assumed individuals will make a guess in some answer—like exams or quiz shows—we are often willing to ascribe someone knowledge merely on the basis that they are capable of supplying a correct answer (Woozley (1953, 155), Hawthorne (2003, fn53), cf. Radford (1966)). To settle which questions we are willing to say an exam candidate or quiz show contestant knew the answers to, we typically need to do no more than check which questions they answered correctly.

We can strengthen this observation by considering the role that denials of knowledge can play in explanation in the context of exams or quizzes. Whereas (5.a) provides an adequate explanation of getting an answer wrong, (5.b) does not.

(5) I got only $\frac{9}{10}$ on the exam because...
   a. ...I didn't know the correct answer to Q2.
   b. ...I wasn't confident about the correct answer to Q2.

On views which take knowledge to be strong, this contrast is surprising. After all, if knowledge is strong then failing to know the correct answer does not preclude being able to supply it on demand (e.g., as a guess) any more than failing to be confident does. (Indeed, if knowledge required a doxastic commitment at least as strong as being confident, then failing to be confident would entail failing to know.) In contrast, if knowledge is weak, then as long as it is common ground that individuals will guess whichever complete answer is most likely on their evidence (and aren't gettierized), denying someone knew will entail that the guess they made was incorrect. Accordingly, it is unsurprising that (5.a) would explain the failure to get full marks. What is distinctive about exams and quiz shows is that the cost of making an incorrect guess is the same as the cost of failing to make a guess at all. So, they are precisely the type of situation in which it will be common ground that individuals will make a guess in some complete answer.

Second, question-answer pairs involving constituent (wh-) questions provide further evidence for the weakness of knowledge. Where a group of individuals each made guesses about the election outcome based on the polling data, (6) can be used to inquire about what they guessed.

(6) Who knew that it would be A who won?

It is widely assumed that constituent questions—like (6)—carry a form of existential presupposition (Katz (1972); Keenan & Hull (1973); Comorovski (1996); Dayal (1996)). For example, asking 'Who broke the vase?' is odd if it is not common ground (and cannot be easily accommodated) that the vase was broken by someone. If good guessing were insufficient to meet the doxastic requirements on knowledge, we would expect (6) to be marked in context (due to failure of the existential presupposition). Yet, it is perfectly felicitous. This can be explained only if presupposing that a group of agents (merely) guessed the answer to a question is compatible with presupposing that at least one of them knows.

Third, the account is supported by judgments about discourse relevance. One test for discourse relevance is provided by licensing of tag clauses in responses to polar question. 'yes'/'no'-answers allow for elucidation, whereby they are accompanied by an supplementary information, as in (7):

(7) a: Did anyone know who would win the election?
b: i. Yes, Lia guessed that it would be A.
ii. ?? Yes, Lia hoped that it would be A.

On views which take knowledge to be strong, the contrast between the i- and ii-response is unexpected. In order for an elucidation on a 'yes'/'no'-answer to be felicitous, it must provide support for the answer which it supplements. But, if knowledge requires a demanding form of doxastic commitment, information about what Lia guessed should be no more relevant to what she knew than information about what she hoped. (In fact, insofar as the i-response will carry a scalar implicature that she didn't bear a stronger attitude to the claim that A would win, it would be expected implicate that Lia didn't know who would win).

A proponent of the view that knowledge is strong might wish to acknowledge these examples, but to bracket them as the product of a pragmatic mechanism (such as loose talk). This strategy shouldn't be ruled out entirely. However, it is unclear whether it has much to recommend it. Absent independent evidence to the contrary, our default hypothesis should be that our judgments are accurately tracking the truth-conditions. Accordingly, insofar as **Knowing** readily accommodates these uses, it is at an explanatory advantage.

Furthermore, when we consider the wider evidence, it is in tension with the hypothesis that there is a general pragmatic mechanism by which knowledge ascriptions can sometimes be used to communicate that an agent is in a state with a less demanding doxastic requirement. If there were such a mechanism, we would expect it to apply generally, both to other attitudes with a demanding doxastic requirement and to demanding doxastic attitudes themselves. But this is not what we find.

(8) a. Lia was [confident/sure/convinced] that A would win the election.
b. Lia was [sad/happy/angry/pleased] that A would win the election.

In a context in which Lia merely guessed A would win, it would be inappropriate to report her pre-election mental state using any version of (8.a). Similarly, even if it assumed that she had the relevant affective state, it would be inappropriate to describe her using any version of (8.b).

3.2 Assertion

Whether it is appropriate to ascribe knowledge on the basis of a (mere) guess appears to depend on the information of the ascriber. Imagine, prior to the election, asserting one of (9) or (10) (when the only evidence you possess about the outcome is the polling data). Unlike with (3)-(4), there is a strong sense that making either of these assertions would be inappropriate.

(9) Lia knows who will win the election.

(10) Lia knows that it will be A who wins the election.

In fact, we can observe something stronger. Based on the polling data, it would seem reasonable for someone to assert (11) prior to the election.

(11) Lia does not know who will win the election.

At first glance, this looks like problem for the view that knowledge is weak. After all, if A is in fact going to win the election, then the former sentences are predicted to be true and and the latter to be false.

There is, however, a readily available explanation. A wide range of authors have argued that assertion is subject to a knowledge norm. Although details differ, the idea is—roughly—that you must not assert what you do not know (Unger (1975); Williamson (1996, 2000); Adler (2002); DeRose (2002); Hawthorne (2003); Schaffer (2008); Engel (2008); Turri (2010, 2011, 2015, 2016)).

Yet, if assertion is subject to a knowledge norm, then this kind of sensitivity in knowledge ascriptions to ascribers' evidence is to be expected.

On the present account, knowledge is a three-place relation between an agent, a proposition and a question. Accordingly, there is an issue of how exactly to formulate the knowledge norm. Schaffer (2008) argues for the following constrastivist implementation:

K-Norm You must: assert p in a context c only if you know p relative to Q and Q is salient in c.

This formulation of the norm agrees with much work in linguistics in holding that what it is permissible to assert depends on what question is salient. The idea is that each (permissible) assertion in a conversation is aimed at resolving a salient question (often identified with the question under discussion or QUD; Roberts (1996, 2012); Ginzburg (1996)). The contrastivist version of the knowledge norm adds that, to be permissible, the speaker must know the proposition asserted relative to the question it aims to resolve.

When combined with **K-Norm**, the view has the resources to explain the judgments above. Assume—as seems reasonable—that an assertion of (9) aims at resolving the polar question, *Does Lia know who will win?*.[9] Then (9) is predicted to be unassertable for anyone who, like Lia, matches their credences to pre-election polling.

Why is that? On any rational way of assigning credences, *Lia knows who will win* can be no more probable than *A will win*. For the former to be true, Lia's guess about *Who will win?* must be a cogent complete answer and that answer must be true. But the only cogent complete answer to the question is *A will win*. According to the polling data, it is more likely that A will lose than that A will win. So, for anyone who bases their credences on the polling data, *Lia knows who will win* will not be a good guess about the question *Does Lia know who will win?*. And since it isn't a good guess about that question, it can't be known relative to it either. So, by the knowledge norm, someone who, like Lia, matches their credences to pre-election polling will not be in a position to assert (9). By the same reasoning, asserting (10) is also predicted to be illicit, assuming that the (unique) salient question which it answers is *Who will win?*.[10]

A further step is needed to explain why (11) appears reasonable to assert—but it is only a small step. Whether an action satisfies its primary norms is not always all that matters for evaluating it. Often, it also matters how it does so. For example, someone who makes a promise they do not expect to keep will be liable to criticism even if, contrary to their expectation, they do what they

[9] If, instead, the assertion aims at resolving a non-polar question, such as *Who knows who will win?*, the dialectic is largely unchanged. Based on the pre-election polling data, the best guess about this question is that no-one knows who won (i.e., the denial of the existential presupposition).

[10] If the polar question *Will A win?* is salient, (10) will be straightforwardly unassertable due to being false—*A will win* is not a good guess about this question, given the polling data.

promised to do. In contrast, someone who expects to do what they promise but is unable to due to unexpected circumstances may be excused, even if, in failing to do as promised, they are acting impermissibly.

Following Williamson (forthcoming) and Littlejohn (forthcomingb,f), I'll describe norms specifying how other norms ought to be satisfied as secondary norms. For instance, it seems reasonable to think there is a norm requiring agents not to act in ways that they do not take to be permissible. Someone who doesn't take themselves to be doing what is required of them may be liable to criticism even if, in fact, they are acting permissibly. Correspondingly, someone who takes themselves to be doing everything required of them may be excused, even if, in fact, they are acting impermissibly.

Depending on the strength of doxastic commitment involved, this secondary norm could take different forms. However, I suggest that we should, minimally, accept the following, where N_{X} is a norm governing X-ing.

G-Norm You must: X only if you guess that you satisfy N_{X} (relative to the question of whether you satisfy N_{X}).

G-Norm says that you mustn't do something unless you are happy to guess that what you are doing is permissible. Taken together, **K-Norm** and **G-Norm** suffice to explain why ascribing Lia ignorance seems reasonable.

Given that it is more likely that A will lose than that A will win, *Lia does not know who will win* is a good guess about the polar question *Does Lia know who will win?*. And, correspondingly, the speaker's best guess about whether they know this proposition (relative to the polar question) is that they do.[11] That's because it is more likely that their best guess is true than that it is not. But, if your guess about whether your assertion satisfies its norm is that it does, then making that assertion will satisfy the requirements of **G-Norm** (even if it is not, in fact, permissible). So, insofar as satisfying **G-Norm** provides an excuse in cases where you fail to satisfy a primary norm, someone who denies that Lia knows prior to the election will be excused.

Combining **Knowing** with an appropriate package of primary and secondary norms explains the sensitivity of our judgments about knowledge ascriptions to variation in ascriber's evidence.[12] That's because, unlike on views which take knowledge to be strong, our judgments about (9)-(10) and (11) can be attributed to facts about their assertability (which is sensitive to the speaker's evidence) rather than truth-value (which isn't). This diagnosis is supported by a pair of further data points.

[11] Assuming, that is, that they are certain that they are not gettierized.

[12] To see that it is asymmetries in evidence, rather than tense, which explains the contrast between (3)-(4) and (9)-(11), imagine that there is a (well-executed) plot to fix the election which will ensure that A wins. Imagine, however, that Lia (who is unaware of the plot) matches her credences to the polling data, as before. In this variant case, it would not be illicit for someone aware of the plot to assert either of (9) or (10). On the contrary, both would be natural ways for someone to communicate to an audience which is also aware of the plot that Lia has correctly guessed the winning candidate.

First, it provides a better fit with observations to do with retraction. If A is revealed to be the winner, someone who previously asserted that Lia did not know who would win will be under pressure to retract what they said (by saying, e.g., 'I was wrong'/'I take it back'/'Lia *did* know, after all').[13] This pressure to retract is strong evidence that (11) is correctly diagnosed as excusably assertable—though false—rather than true. If (11) were true (as predicted by views which treat knowledge to be strong), it is hard to explain why retraction would be permissible. After all, we are not typically expected to retract true assertions. In contrast, where someone becomes aware that they acted impermissibly despite having an excuse, they typically ought to take steps to rectify their wrongdoing. And where the original act is an assertion, rectifying wrongdoing will typically require retraction.

Second, it can explain some otherwise puzzling judgments about embedded knowledge ascriptions raised in Holguín (2021). It would be acceptable for someone (prior to the election, who matches their credence to the polls) to assert any of (12)-(14):

(12) If the polls are right, then Lia knows who will win.

(13) Either A is going to lose or Lia knows who is going to win.

(14) Lia might know who the winner is.

On views which take knowledge to be strong—and, so, treat (9) as false—this is, again, hard to explain. On such views, each of the three complex sentences above should be predicted to be false as well. In contrast, on views which treat (9) as true but unassertable, this behavior is easy to explain. If knowledge is weak, then (12)-(14) are predicted to be true irrespective of who wins the election. Accordingly, someone who matches their credence to the polls ought to assign maximal credence to each of the propositions expressed by the sentences. It follows that not only will each of these propositions be knowable (relative to any choice of question), it will be a good guess about whether it is knowable that it is. So, unlike (9) and (10), assertions of (12)-(14) are predicted to satisfy both **K-Norm** and **G-Norm**, and, thus, to be assertable.

4 Objections (& Replies)

In this section I consider four types of objection to the account developed above. These objections draw on observations to do with denials of knowledge (§4.1); lotteries (§4.2); Moorean assertion (§4.3); and closure principles (§4.4). I argue that **Knowing** has the resources to adequately address each of them.

[13] Retraction need not be obligatory. The fact that it is permissible is, equally, evidence for the weakness of knowledge.

4.1 Ignorance

Objection: If knowledge is weak, denials of knowledge should be correspondingly strong. Yet this is not what we observe.

Reply: Ascriptions of ignorance are often acceptable in circumstances in which agents would be predicted to possess knowledge, according to **Knowing**. First, overt modification can lower the standards for ascribing ignorance. For example, focal stress (as in (15.a)) and lexical intensifiers (as in (15.b)) can make denying Lia knowledge that A won acceptable, even if her guess about who would win has been established to be correct.

(15) a. Lia didn't *know* it would be A who won.
b. Lia didn't [really know/know for sure] it would be A who won.

Second, in the appropriate contexts, ignorance ascriptions of this kind can be licensed even in the absence of overt modification. Utterances like (16.a), simultaneously self-ascribing ignorance and a correct guess, can be acceptably asserted without overt modification of 'knows'. The same goes for (16.b), which is assertable even in circumstances where many non-compatriots have made a (good) guess that their country's team will lose (and only one team will win).

(16) a. I didn't know whether it would rain, although I correctly guessed that it would.
b. Nobody knows that their country's team won't win the world cup (though many people will have correctly guessed it won't).

These examples constitute important data which need to be explained. However, on reflection, they cannot serve as the basis of a stable argument against the claim that knowledge is weak. That's because ignorance ascriptions of the same kind are also acceptable in cases where an agent has extremely high (rational) credence. Accordingly, to the extent that such examples count as evidence that knoweldge is not weak, they will also count as evidence that it is implausibly strong.

First, overt modification (either via focus or lexical intensifiers) can make ignorance ascriptions acceptable for even overwhelmingly well-supported propositions—as in (17.a-c) (cf. Unger (1975, §II.7-9), Ginet (1975, §2)).

(17) Lia doesn't [*know*/really know/know for sure] that
a. ...the bank will be open.
b. ...the sun will rise tomorrow.
c. ...she's not a BIV.

And, second, utterances like (18.a-b) involving a doxastic state stronger than guessing but weaker than certainty are also clearly acceptable in the absence of modifiers.

(18) a. I didn't know whether it would rain, although I was [confident/pretty sure/almost certain/?certain] that it would.

b. Nobody knows that their country's team won't win the world cup (though many people are [confident/pretty sure/almost certain/?certain] it won't).

Based on examples like these, it seems clear that 'knows' sometimes expresses a state requiring an extremely strong form of doxastic commitment (i.e., at least entailing certainty). What is important, in assessing the present objection, is whether it always expresses such a state.

Both those who take knowledge to be weak and those who take it to be strong without requiring certainty ought to deny that it does. They should agree that, in typical uses, 'know' expresses a state with a less demanding doxastic requirement. It is this state which—they will claim—is the object of ordinary, non-skeptical theorizing in epistemology, rather than the stronger state expressed when, e.g., 'knows' occurs under focus (Stanley (2008)). And, although they disagree over whether that state's doxastic requirement is weak or (merely) moderately strong, evidence provided by (15)-(18) will not help to settle this question.[14]

This point is not novel. It is a relatively common observation that knowledge ascriptions (whether overtly modified or not) can sometimes express an extremely demanding state. However, this state is not usually taken to be of primary interest to epistemologists. Rather, it is standard to focus instead on the less demanding state which, it is claimed, is expressed in most ordinary knowledge ascriptions. There are various possible views on how 'knows' comes to express a demanding attitude in some uses (whether via context-sensitivity (Cohen (1986, 1987, 1988); DeRose (1992, 1995, 2000); Lewis (1996)), polysemy/ambiguity (Malcolm (1952); Feldman (1986, §2); Steup (2005); Satta (2018a,b,c, 2020)) or something else). In the present setting, however, the precise mechanism is not crucial to the dialectic. Once we grant that 'knows' does not typically express a state entailing certainty, the examples above give us no further reason to deny that in many ordinary uses it expresses a state which is weak.[15]

[14]Some authors have defended the view that (true) knowledge ascriptions do require certainty (Moore (1959); Miller (1978); Klein (1981)). To avoid skepticism, however, they take certainty to be (relatively) undemanding—in the relevant sense, we are often certain that, e.g., the sun will rise tomorrow/the bank is open/we are not BIVs (cf.Beddor (2020); Goodman & Holguín (forthcoming); Unger (1975) is an exception). Accordingly, (17.a-c) will equally put pressure on these authors to insist that modified knowledge ascriptions express a state which requires a commitment stronger than certainty.

[15] Can 'knows' ever express a state which requires a doxastic commitment intermediate between guessing and certainty? There is some evidence that it can. Imagine that Xena, Yuli and Zara differ in their credence that A won. Xena takes it to be overwhelmingly likely; Yuli merely takes it to be more likely than not; finally, Zara takes it to be less likely than not (though she still thinks that A is the most likely winner).

(†) Only Xena and Yuli know who won.

4.2 Lottery Knowledge

Objection: Imagine that Ali has one ticket in a 10,000-ticket fair lottery. According to orthodoxy, knowledge ascriptions like (19.a-b) are false in all contexts.

(19) a. Ali knows whether his ticket will win.
b. Ali knows that his ticket will lose.

Yet, that Ali's ticket will lose is a good guess about whether it will win. So, assuming Ali does in fact have a losing ticket and guesses that he does, (19.a-b) are predicted to be true in any context where the polar question is uniquely salient.[16]

Reply: To start with, it is worth briefly considering some evidence that—contrary to orthodoxy—(19.a-b) are (at least sometimes) true in the kind of context described (cf. Hawthorne (2003, 18-20), Sosa (2015)).

First, observe that the acceptability of knowledge ascriptions in lottery cases is supported by the same evidence as knowledge ascriptions in our original case of weak knowledge (§3.1). It would be reasonable for someone aware of his guess to ask 'How did Ali know his ticket would lose?' (to which the obvious answer is 'Because it was so unlikely to win'). The same goes for evidence involving constituent questions and discourse relevance.

Second, in many settings it appears natural to ascribe Ali awareness that his ticket will not win (based on the low chances alone). Imagine that Ali is behaving overly-optimistically about his chance of success (making holiday plans, putting champagne on ice, etc.). It would be unobjectionable for someone to, rhetorically, utter (20.a). And, equally, it would be reasonable for Ali to respond by saying something along the lines of (20.b).

(20) a. Ali is aware he's not going to win, right?
b. Look, of course I'm aware I'm not going to win—I am just having fun imagining it.

But (propositional) awareness entails knowledge; anyone aware that p, knows whether p (cf. Williamson (2000, §2)). So, if examples like this demonstrate we

Arguably, (†) has a true reading in these circumstances. If that's right, one option would be to attribute this to context-sensitivity in the state expressed by 'knows'. In some, but not all, contexts it expresses a state which requires doxastic commitment stronger than guessing but weaker than certainty (for example, Goodman & Salow (2021) defend a version of this view). Another option would be to opt for a more restrictive account of good guessing, for example, by replacing **Guessing** with **Guessing*** (fn.3). On the latter, wherever all three individuals adopt an epistemic utility function which values accuracy sufficiently highly relative to informativity, the latter will predict that only Xena and Yuli will be in a position to make a good guess in a complete answer to *Who won?*. What is crucial for present purposes, however, is that (†) also has a reading on which it is false. To explain that reading, 'knows' must also be capable of expressing a state which has a weak doxastic requirement.

[16] To control for worries about gettierization, we can stipulate for the sake of the objection that (unbeknownst to Ali) the lottery has been rigged so that Ali couldn't easily have lost.

can attribute Ali awareness that he will lose, they also demonstrate that we can attribute him knowledge.

Third, there is a contrast in sentences like (21) between ascriptions of knowledge and ascriptions of emotive attitudes (like regret or disappointment).

(21) Ali [knows/??regrets/??is disappointed] that the ticket he bought will lose.

It seems notably worse to describe Ali as bearing an attitude such as regret or disappointment to the fact his ticket will lose than to describe him as knowing it. An obvious explanation of this contrast, if knowledge is weak, is that the form of doxastic commitment required by attitudes such as regret/disappointment is more demanding than that required knowledge. While guessing that your ticket will lose suffices for the latter, it does not suffice for the former.

I have suggested that there are grounds for hesitancy regarding the judgment that it can never be known that a lottery ticket will lose on statistical evidence alone. With that said, it is important to consider the objection's force—or lack of it—in relation to the broader question of the strength of doxastic commitment required for knowledge.

Even if (19.a-b) are false, **Knowing** is still no worse off than other anti-skeptical accounts. Among those who deny knowledge based on statistical evidence in lotteries, it is a point of agreement that the size of the lottery does not make a difference. Accordingly, any view on which knowledge does not require certainty will need to say something more to distinguish Ali's case from cases of genuine knowledge—the strength of doxastic commitment cannot be what explains the unavailability of knowledge. In this respect, lotteries do not pose a special problem for views on which knowledge is weak. It remains unclear what this distinguishing feature might be (cf. Hawthorne (2003); Pritchard (2008) for discussion). However, there is no reason to think that whatever is distinctive about lotteries (if, in fact, anything is) cannot equally be cited by the proponent of **Knowing**.

Those who deny that you can know you will lose a lottery on the basis of merely statistical evidence will presumably take the same position in cases where the likelihood of error is higher. They will also deny you could—for example—know how a flipped coin will land merely on the basis that it has a slight bias towards heads. In contrast, those who take knowledge to be weak and to be available in the lottery may well wish to allow for knowledge cases like (slightly biased) coin flips. This prediction is not forced, however. The two cases could come apart, for example, in the presence of principles connecting knowledge and objective chance. If sufficiently high (past or present) objective chance of error suffices for gettierization, then the high pbjective chance of tails may act as a barrier to knowledge even if the low chance of winning the lottery doesn't (cf. Hawthorne & Lasonen-Aarnio (2009); Dorr *et al.* (2014)). Insofar as high objective chance is compatible with arbitrarily low rational credence, the claim that knowledge is weak is silent on such principles. As a result, it is also silent on whether the availability of lottery knowledge implies the availability of knowledge of the outcome of other chancy processes.

4.3 Moorean Conjunctions

Objection: It is widely recognized that sentences like (22)—so-called 'Moorean conjunctions'—are unassertable in any context.

(22) I don't know where my keys are but they're on the table.

However, according to **Knowing**, there are questions relative to which a speaker can know the proposition expressed by her own utterance of (22). Doesn't this show that some aspect of the view must be wrong?

Reply: It's correct that, where p is an answer to Q, there exist questions relative to which S can know the proposition *p and there is no $q \in Q$ such that S knows q relative to Q*. However, knowledge does not suffice for assertability on view above. Permissible assertion requires not only that the proposition asserted is known relative to the contextually salient question(s) (by **K-Norm**)—it also requires that it is a good guess, for the speaker, that it it is known (by **G-Norm**). And, crucially, this latter condition is unsatisfiable for Moorean conjunctions.

First, observe that sentences like (23)—where the right-hand conjunct is a merely partial answer to the embedded question in the left-hand conjunct—are unobjectionable to assert.

(23) I don't know where my keys are but they aren't in my pocket.

Accordingly, we can restrict our attention to the case where the right-hand conjunct is a complete answer, instead. Let p' be the proposition *p and there is no $q \in Q$ such that S knows q relative to Q*. We'll show that wherever p is a complete answer to Q (i.e., $p \in Q$) and p is S's guess about Q, there is no question relative to which S can permissibly assert p'.

Suppose, first, that p is a cogent answer to Q. Then, every world at which p is true is a world at which S makes a correct good guess about a complete answer to Q. So the only worlds at which p' is true are worlds at which S is gettierized about p. But, if S is gettierized about p, then S is plausibly also gettierized about $p \wedge q$ (for any q).[17] So, at any world at which S is gettierized about p, S will also be gettierized about p'. Now consider any Q' to which p' is an answer (partial or complete). It follows that there are no worlds at which S can know p' relative to Q' (since at every world, either p' is false or S is gettierized about p') and so their credence that they know p' relative to Q' should be 0. Accordingly, S's best guess about *Does S know p' relative to Q'?* will be that they do not.

[17] Not only is this motivated by judgments about individual cases, it follows from a treatment of non-gettierization in terms of reliability (of the kind adopted in §5).

Suppose, instead, that p is not a cogent answer to Q. Then the credence S assigns to p must be below .5 and, likewise, so must the credence they assigns to p' (since p' entails p). So consider, again, any Q' to which p' is an answer. Since S's credence in *S knows* p' *relative to* Q' can be no greater than S's credence in p', it must be below .5, too. Accordingly, S's best guess about *Does S knows* p' *relative to* Q'*?* will be that they do not.

By **G-Norm**, you must not assert a proposition if your best guess about whether you know it relative to the salient question(s) is that you don't. But, in the argument above, Q' was arbitrary. So there is no question relative to which it will be permissible for S to assert p'.[18]

4.4 Closure

Objection: Knowledge is reliably transmitted via deductive inference. But **Knowing** predicts failures of both multi- and single-premise closure. So it can't explain the role of deductive inference in transmitting knowledge.

Reply: There are a variety of different single-premise closure principles which can be formulated in a contrastivist setting (cf. Schaffer (2004, 2007a); Hughes (2013); Kvanvig (2007, 2008); Hawke (2016)). It is important to distinguish between these principles in assessing what a contrastivist account predicts about knowledge transmission.

The present account predicts that knowledge is closed under both multi- and single-premise closure relative to a fixed question (as long as attention is restricted to answers to that question).

Closure$_1$ Suppose that S knows p relative to Q and S knows q relative to Q. Then S knows $p \wedge q$ relative to Q.

Closure$_2$ Suppose that S knows p relative to Q. Then for any q entailed by p which is an answer to Q, S knows q relative to Q.

In contrast, knowledge is not closed under either multi- or single-premise closure relative to different questions.

Closure$_3$ Suppose S knows p relative to Q and S knows q relative to Q'. Then, if $p \wedge q$ is an answer to Q'', S knows $p \wedge q$ relative to Q''.

Closure$_4$ Suppose that S knows p relative to Q. Then for any q entailed by p which is an answer to Q', S knows q relative to Q'.

Closure$_3$ and **Closure$_4$** fail for related reasons. Cogency is not, in general, preserved across questions. Against **Closure$_3$**, observe that from the fact that

[18] The same reasoning generalizes to explain the unassertability of sentences like (‡), on the assumption that by asserting (‡), S will raise to salience some question to which *S's keys are on the table* is a complete answer.

(‡) I don't know that my keys are on the table but they are on the table.

each of a pair of answers is entailed by a cogent answer to one of a pair of questions, it does not follow that their conjunction will be entailed by a cogent answer to a third question. And against **Closure**$_4$, observe that from the fact that an answer is entailed by a cogent answer to one question, it does not follow that its consequences will be entailed by a cogent answer to a second.[19]

In both the multi- and single-premise cases, failure of the more general closure principles is motivated. Failures of **Closure**$_3$ are supported by considering what can be known in cases with a preface-like structure (Makinson (1965); see Littlejohn & Dutant (2020); Carter & Goldstein (2021, 2023); Goodman & Salow (2023); Dutant & Littlejohn (forthcoming) for discussion of knowledge and the preface). Imagine a historian has written a book comprising 1,000 claims about Napoleon I. Each of the claims in her book is well-researched. However, upon submitting the manuscript, she receives a message from an extremely reliable fact-checking service reporting that (exactly) one claim in her book has been found to be false.

(24) The historian knows whether nth claim in the book is true.

(25) The historian knows how many claims in the book are false.

(26) The historian knows which claims (if any) in the book are false.

Assume that the report from the fact-checking service is correct. Given that each claim in her book is well-researched, it seems overwhelming plausible that, for each $1 \leq n \leq 1,000$ such that the nth claim is, in fact, true, you can assert the corresponding instance of (24). Given reliability of the fact-checking service, it also appears plausible that you can assert (25). Yet, equally, on the assumption that her evidence does not favor any individual claims more strongly than others, it also seems plausible you are not in a position to assert (26). Vindicating this pattern of judgments, however, requires giving up **Closure**$_3$.

Failure of **Closure**$_4$ is also motivated, conditional on knowledge being weak. Imagine a deck comprising the following cards: 6♠7♠8♠9♠10♠A♣A♥A♦. Imagine that, after guessing a complete answer to *What suit will be drawn?* and *What value will be drawn?*, you proceed to draw 7♠ from the deck. If knowledge is indeed weak then, assuming your guess was good, you will be in a position to truthfully assert (27). In contrast, you do not appear to be in a position to truthfully assert (28), regardless of what guesses you made about the latter question.

(27) I knew that a spade would be drawn.

(28) I knew that an ace would not be drawn.

[19] Both will also fail for the less interesting reason that an agent may make a guess about one question while failing to make a guess about another. The failure of cogency to be preserved across questions, however, also tells against variants of **Closure**$_{3-4}$ stated in terms of what an agent is in a position to know relative to a question.

The present proposal can predict this pattern of judgments. Assume (as seems natural) that asserting (27) would make *What suit will be drawn?* salient, whereas asserting (28) would make *What value will be drawn?* salient, instead. Then whereas an assertion of (27) would be true, an assertion of (28) would be false: no good guess about the latter question entails that an ace won't be drawn. Crucially, however, this requires a failure of $\textbf{Closure}_4$, since drawing a spade entails not drawing an ace.[20]

There are limits on how failures of $\textbf{Closure}_{3-4}$ can be reported, however. In particular, the view developed above predicts that conjunctions such as (29) and (30) can never be used to assert a truth.

(29) ?? Lia knew who would win but she didn't know whether it would be A.

(30) ?? I knew that a spade would be drawn but I didn't know that an ace would not be drawn.

By assumption, the left- and right-hand conjuncts of (30) raise to salience the questions *What suit will be drawn?* and *What value will be drawn?*, respectively. *A spade will be drawn* is a (partial) answer to both of these questions (since it is equivalent to *An ace will not be drawn*). Accordingly, for the left-hand conjunct to be true, the speaker must know this proposition relative to both question. But *A spade will be drawn* is not a cogent guess about *What value will be drawn?*. So the left-hand conjunct will be false in any context in which (30) is asserted.

Turning to (29), the left- and right-hand conjuncts of raise to salience the questions expressed by their respective embedded interrogatives. But the question *Who will win?* has *Will A win?* as a part (since every answer to the latter is a partial answer to the former). Accordingly, for the left-hand conjunct of (29) to be true in context, Lia needs to know a complete answer to both questions. Yet this is incompatible with the truth of the right-hand conjunct. So the conjunction will be false.

5 Higher-Order Knowledge

In this section, we'll consider what **Knowing** has to say about the availability of higher-order knowledge. To do this, it will be helpful to formulate the theory within a simple model.

Let W be a domain of worlds. Propositions $p, q, ... \subseteq W$ correspond to sets of worlds. For any proposition p, $\overline{p}$ is the complement of p in W. Questions $Q, Q'...$

[20] Those concerned that the entailments between drawing a spade and not drawing an ace hold only relative to background information about the composition of the deck should feel free to modify the examples to involve claims about what suit/value would be drawn from the specific deck in question (rigidly designated).

correspond to sets of mutually exclusive and jointly exhaustive propositions.[21] A locality relation, R, is a reflexive, symmetric relation on W, representing what is easily possible (relative to some privileged agent represented by the model). $R(w)$ corresponds to the strongest proposition which could not easily have been false at w. Cr is a classical probability measure over the set of propositions, corresponding to a rational way of assigning credences to propositions. As an idealization, we ignore the possibility of uncertainty about the credence function. A model $\langle W, R, Cr\rangle$ is a triple comprising a domain, locality relation, and probability measure. Relative to a model, Cg maps questions to their cogent answers.

Definition 1. $p \in Cg(Q)$ iff there is some non-empty $X \subseteq Q$ such that: (i) $p = \bigcup X$; and (ii) For all $q \in X$ and $r \in Q - X$: $Cr(q) > Cr(r)$.

What can be known is represented by an operator, K, mapping propositions and questions to propositions. $K(p, Q)$ is the proposition that p is a candidate for being the strongest answer known relative to Q.

Definition 2. $K(p, Q) = \{w \mid p \in Cg(Q) \wedge R(w) \subseteq p\}$

That is: $K(p, Q)$ is true iff p is a cogent answer to Q and p couldn't easily have been false. The requirement that what is knowable couldn't easily have been false is intended to model (a consequence of) the non-gettierization condition. The idea is that an agent reliably makes a correct guess only if, given the way they assign credences to answers, they could not easily have made an incorrect guess in the same answer. (Unger (1968); Engel (1992); Zagzebski (1994); Pritchard (2003, 2005); Williamson (2013)).

This implementation involves a level of simplification. There are at least two different ways that an agent could—intuitively—fail to reliably make correct guesses. First, if the agent could easily have rationally assigned the same credences when the underlying state of the world was significantly different (e.g., the polling method could easily have produced the same data based on significantly different voting intentions). Second, if the agent could easily have rationally assigned significantly different credences while the underlying state of the world was the same (e.g., the polling method could easily have produced significantly different data based on the same voting intentions). The present model captures only the former type of unreliability. As a result, it fails to accommodate the possibility of unreliable guesses in modally robust truths (e.g., nomological necessities, etc.). However, this simplification is harmless. If the results below hold for agents who could not have easily rationally adopted different credences, they will hold more generally.

As discussed above (§3.2), knowledge does not iterate freely on the present proposal. You can know the answer to a question while being unable to know the answer to whether you can know the answer to that question. There are

[21] That is, $Q \subseteq \mathcal{P}(W)$ is a question iff (i) $\bigcup Q = W$ and (ii) for all $p, p' \in Q : p \cap p' = \varnothing$.

many views on which knowledge fails to iterate. However, the way such failures arise in the present framework is distinctive.

To state principles about how knowledge iterates within the framework, we first need to be able to state questions about what can be known at an arbitrary order. Observe that $\bigcup_{p \in Q} K(p, Q)$ is the proposition that some complete answer to Q can be known relative to Q. So, for a given choice of question, Q, we can recursively define a sequence of questions as follows:

- $Q_0 = Q$;
- $Q_{n+1} = \{\bigcup_{p \in Q_n} K(p, Q_n), \overline{\bigcup_{p \in Q_n} K(p, Q_n)}\}$

Informally, Q_n is the question *Is any complete answer to Q_{n-1} knowable relative to Q_{n-1}?* (where $n \geq 1$). So, $Q_0, Q_1, ...$ is a sequence of questions starting with Q and where each subsequent question asks whether some complete answer to the preceding question can be known relative to itself. Let **Introspection** be the principle that being able to know a complete answer to Q_n implies being able to know a complete answer to Q_{n+1}.

Introspection For all n: $\bigcup_{p \in Q_n} K(p, Q_n) \subseteq \bigcup_{p' \in Q_{n+1}} K(p', Q_{n+1})$

Failures of introspection principles are commonly motivated by a margin-for-error constraint (Williamson (2000, 2011, 2013, 2014)). This constraint is, in turn, motivated by the idea that what an agent knows could not easily have been false. What is distinctive about the present framework is that **Introspection** fails for reasons which have nothing to do with what could easily have been the case.

Say that a locality relation, R, is trivial iff $R = \{\langle w, w \rangle \mid w \in W\}$. In models with a trivial locality relation, nothing true could easily have been false.

Observation 1. There exist models with a trivial locality relation in which **Introspection** fails.

In models with a trivial locality relation, the requirement that what is knowable could not easily be false collapses to the requirement that what is knowable is true. Even if we restrict our attention to such models, it will not in general be the case that if a complete answer to Q_n can be known (relative to Q_n) then it can be known (relative to Q_{n+1}) that a complete answer to Q_n can be known (relative to Q_n). Or, more informally, knowing a complete answer to a question does not entail knowing a complete answer to whether you know a complete answer to that question.

To see why, consider the four-world model depicted in Figure 1. Assume that each world is locality-related only to itself. The blue partition corresponds to the question $Q_0^* = \{\{\mathsf{A}\}, \{\mathsf{B}\}, \{\mathsf{C}\}, \{\mathsf{D}\}\}$. Its unique cogent complete answer is

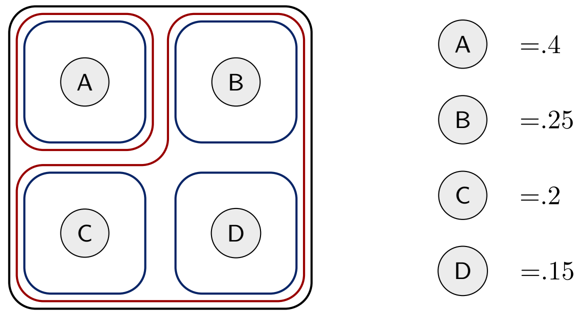

Figure 1: A four-world model.

{A}. At each world, all and only the true cogent answers to a question are known. Accordingly, the only world at which a complete answer to Q_0^* is known relative to Q_0^* is A. The red partition corresponds to the question Q_1^*. Its unique cogent complete answer is $\overline{\bigcup_{p \in Q_0^*} K(p, Q_0^*)}$ (i.e., the proposition that no complete answer to Q_0^* can be known (relative to Q_0^*)). Accordingly, the only worlds at which a complete answer to Q_1^* can be known relative to Q_1^* are B, C, and D. So, although there are worlds at which a complete answer to Q_0^* can be known (relative to Q_0^*), there is no world at which this is can be known (relative to Q_1^*). Hence, **Introspection** fails.

While **Introspection** can fail in models with a trivial locality relation, its failures are limited. In particular, **Introspection** fails only for sequences whose initial question has no complete answer which is more likely than its negation. If there is some $p \in Q_0$ such that $Cr(p) > .5$, then for all n: $\bigcup_{q \in Q_n} K(q, Q_n)$ implies $\bigcup_{q' \in Q_{n+1}} K(q', Q_{n+1})$.[22] However, this is an artifact of the restriction to models in which the locality relation is trivial. When we consider the full range of models, we find that **Introspection** can fail in any sequence, as long as no answer to the initial question has maximal credence.

More generally, there is no non-maximal credence in p which suffices to guarantee that *S knows p relative to Q* will be a good guess about *Does S know p relative to Q?*.

Observation 2. There is no $t < 1$ such that in all models: if $p \in Q$ and $Cr(p) \geq t$, then $K(p, Q) \in Cg(\{K(p, Q), \overline{K(p, Q)}\})$.

That's because as long as $p \subset W$, the credence assigned to the proposition that p could easily have been false (i.e., $\{w \mid \overline{p} \cap R(w) \neq \varnothing\}$) can be arbitrarily high. And where p could easily have been false, p cannot be known relative to any question. So, there will be no $t < 1$ for which it is guaranteed that if

[22] As a corollary, sufficiently high-order knowledge does iterate freely in such models. Specifically, $\bigcup_{p \in Q_n} K(p, Q_n)$ implies $\bigcup_{p' \in Q_{n+1}} K(p', Q_{n+1})$ as long as $n \geq 2$.

$p \in Q_n$ and $Cr(p) > t$, then $K(p, Q_n)$ is a cogent answer to Q_{n+1}.This property of the framework is important. Given the package of norms governing assertion proposed in §3.2, it implies that there is no non-maximal credence in the correct answer to a question which suffices to make it permissible to take yourself to know that answer. And so, correspondingly, there is no non-maximal credence in the correct answer to a question which suffices to make it permissible to assert that answer (in contexts where that question is under discussion).

6 Conclusion

Knowledge is weak. To know the answer to a question, it's enough that you make a good guess about it. And for your guess about a question to be good, it's enough that your rational credence in it is non-zero. Obviously, knowing requires more to go right than simply good guessing. And good guessing requires more than simply non-zero credence. But crucially, there is no positive threshold an agent's credence in an answer must exceed to satisfy the doxastic constraint.

Holguín (2022) and Dorst & Mandelkern (2023) argue for an account of belief in terms of good guessing.

Believing S rationally believes that p (relative to a question Q) iff:
(i) p is an answer to Q;
(ii) S makes a good guess about Q which entails p.

According to **Believing**, making a good guess about a question suffices for having a rational belief about its answer. The arguments above are neutral on the status of **Believing**. However, the combination of **Knowing**, **Guessing**, and **Believing** —which together imply that knowledge is just ungettierized rational true belief—makes for a natural package.

In particular, observe that those who accept **Believing** and **Guessing** face an significant challenge: what role does belief/guessing play over and above credence? Or, put another way: once it is settled what credence an agent assigns to a question's answers, why does it matter which answer they believe/guess? **Knowing** provides a clear answer to this question. Guessing (and, hence, believing) suffices to satisfy the doxastic requirement on knowledge. While what an agent is in a position to know depends on their credences, what an agent knows depends on their guesses. Accordingly, if inquiry aims at knowledge, guessing will play a central role in inquiry. Successful inquirers do not merely need to adopt accurate credences. They also need to make correct guesses.

References

Adler, J. 2002. *Belief's Own Ethics*. Boston, MA: MIT Press.

Beddor, Bob. 2020. New Work For Certainty. *Philosophers' Imprint*, **20**(8).

Carrier, L. S. 1977. The Irreducibility of Knowledge. *Logique Et Analyse*, **77**(Sommaire), 167–176.

Carter, Sam, & Goldstein, Simon. 2021. The Normality of Error. *Philosophical Studies*, **178**(8), 2509–2533.

Carter, Sam, & Goldstein, Simon. 2023. Getting Accurate about Knowledge. *Mind*, **132**(525), 158–191.

Cohen, Stewart. 1986. Knowledge and Context. *Journal of Philosophy*, **83**(10), 574–583.

Cohen, Stewart. 1987. Knowledge, Context, and Social Standards. *Synthese*, **73**(1), 3–26.

Cohen, Stewart. 1988. How to Be a Fallibilist. *Philosophical Perspectives*, **2**, 91–123.

Comorovski, Ileana. 1996. *Interrogative Phrases and the Syntax-Semantics Interface*. Springer Netherlands.

Dayal, Veneeta. 1996. *Locality in WH Quantification*. Studies in Linguistics and Philosophy, vol. 62. Dordrecht: Springer Netherlands.

DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, **52**(4), 913–929.

DeRose, Keith. 1995. Solving the Skeptical Problem. *In:* DeRose, Keith, & Warfield, Ted A. (eds), *Skepticism: A Contemporary Reader*. Oup Usa.

DeRose, Keith. 2000. Now You Know It, Now You Don't. *The Proceedings of the Twentieth World Congress of Philosophy*, **5**, 91–106.

DeRose, Keith. 2002. Assertion, Knowledge, and Context. *Philosophical Review*, **111**(2), 167–203.

Dorr, Cian, Goodman, Jeremy, & Hawthorne, John. 2014. Knowing against the Odds. *Philosophical Studies*, **170**(2), 277–287.

Dorst, Kevin. 2019. Lockeans Maximize Expected Accuracy. *Mind*, **128**(509), 175–211.

Dorst, Kevin, & Mandelkern, Matthew. 2023. Good Guesses. *Philosophy and Phenomenological Research*, **105**(3), 581–618.

Dutant, Julien, & Littlejohn, Clayton. forthcoming. What Is Rational Belief? *Noûs*, **n/a**(n/a).

Engel, Mylan. 1992. Is Epistemic Luck Compatible with Knowledge? *Southern Journal of Philosophy*, **30**(2), 59–75.

Engel, Pascal. 2008. In What Sense Is Knowledge the Norm of Assertion? *Grazer Philosophische Studien*, **77**(1), 45–59.

Feldman, Fred. 1986. *A Cartesian Introduction to Philosophy*. McGraw-Hill Companies.

Feldman, Richard. 2007. Reasonable Religious Disagreements. *Pages 194–214 of:* Antony, Louise (ed), *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. Oxford University Press.

Fodor, Jerry A. 1998. *Concepts: Where Cognitive Science Went Wrong.* Oxford University Press.
Ginet, Carl. 1975. *Knowledge, Perception and Memory.* Dordrecht: Springer Netherlands.
Ginzburg, Jonathan. 1996. Interrogatives: Questions, Facts and Dialogue. *In:* Lappin, Shalom (ed), *The Handbook of Contemporary Semantic Theory.* Blackwell Reference.
Goodman, Jeremy, & Holguín, Ben. forthcoming. Thinking and Being Sure. *Philosophy and Phenomenological Research*, 1–27.
Goodman, Jeremy, & Salow, Bernhard. 2021. Knowledge from Probability. *Electronic Proceedings in Theoretical Computer Science*, **335**(June), 171–186.
Goodman, Jeremy, & Salow, Bernhard. 2023. Epistemology Normalized. *The Philosophical Review*, **132**(1), 89–145.
Groenendijk, Joroen, & Stokhof, Martin. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers.* Ph.D. thesis.
Hamblin, C. L. 1958. Questions. *Australasian Journal of Philosophy*, **36**(3), 159–168.
Hamblin, Charles L. 1973. Questions in Montague English. *Foundations of Language*, **10**(1), 41–53.
Hawke, Peter. 2016. Questions, Topics and Restricted Closure. *Philosophical Studies*, **173**(10), 2759–2784.
Hawthorne, John. 2003. *Knowledge and Lotteries.* Oxford University Press.
Hawthorne, John, & Lasonen-Aarnio, Maria. 2009. Knowledge and Objective Chance. *Pages 92–108 of:* Greenough, Patrick, & Pritchard, Duncan (eds), *Williamson on Knowledge.* Oxford University Press.
Hawthorne, John, Rothschild, Daniel, & Spectre, Levi. 2016. Belief Is Weak. *Philosophical Studies*, **173**(5), 1393–1404.
Hilpinen, Risto. 1988. Knowledge and Conditionals. *Philosophical Perspectives*, **2**, 157–182.
Holguín, Ben. 2021. Knowledge by Constraint. *Philosophical Perspectives*, **35**(1), 1–28.
Holguín, Ben. 2022. Thinking, Guessing, and Believing. *Philosophers' Imprint*, **22**(1), 1–34.
Hughes, Michael. 2013. Problems for Contrastive Closure: Resolved and Regained. *Philosophical Studies*, **163**(3), 577–590.
Karjalainen, Antti, & Morton, Adam. 2003. Contrastive Knowledge. *Philosophical Explorations*, **6**(2), 74–89.
Katz, Jerrold J. 1972. *Semantic Theory.* Harper & Row.
Keenan, Edward L., & Hull, Robert D. 1973. The Logical Presuppositions of Questions and Answers. *Pages 441–466 of:* Petöfi, J'Anos S., & Franck, Dorothea (eds), *Präsuppositionen in Philosophie Und Linguistik - Presuppositions in Philosophy and Linguistics.* Ahtenäum.
Kelly, Thomas. 2013. Evidence Can Be Permissive. *Page 298 of:* Steup, Matthias, & Turri, John (eds), *Contemporary Debates in Epistemology.* Blackwell.
Klein, Peter David. 1981. *Certainty, a Refutation of Scepticism.* University of

Minnesota Press.
Kvanvig, Jonathan L. 2007. Contextualism, Contrastivism, Relevant Alternatives, and Closure. *Philosophical Studies*, **134**(2), 131–140.
Kvanvig, Jonathan L. 2008. Contrastivism and Closure. *Social Epistemology*, **22**(3), 247–256.
Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–567.
Littlejohn, Clayton. forthcominga. Externalism Explained. *In:* Oliveira, Luis R. G. (ed), *Externalism about Knowledge*. Oxford University Press.
Littlejohn, Clayton. forthcomingb. A Plea for Epistemic Excuses. *In:* Julien Dutant, Fabian Dorsch (ed), *The New Evil Demon Problem*. Oxford University Press.
Littlejohn, Clayton, & Dutant, Julien. 2020. Justification, Knowledge, and Normality. *Philosophical Studies*, **177**(6), 1593–1609.
Makinson, David C. 1965. The Paradox of the Preface. *Analysis*, **25**(6), 205.
Malcolm, Norman. 1952. Knowledge and Belief. *Mind*, **61**(242), 178–189.
Miller, Richard W. 1978. Absolute Certainty. *Mind*, **87**(345), 46–65.
Moore, George Edward. 1959. Certainty. *In: Philosophical Papers*. Routledge.
Morton, Adam, & Karjalainen, Antti. 2008. Contrastivity and Indistinguishability. *Social Epistemology*, **22**(3), 271–280.
Pritchard, Duncan. 2003. Virtue Epistemology and Epistemic Luck. *Metaphilosophy*, **34**(1/2), 106–130.
Pritchard, Duncan. 2005. *Epistemic Luck*. Vol. 29. Oxford University Press UK.
Pritchard, Duncan. 2008. Knowledge, Luck and Lotteries. *In:* Hendricks, Vincent (ed), *New Waves in Epistemology*. Palgrave-Macmillan.
Radford, Colin. 1966. Knowledge: By Examples. *Analysis*, **27**(1), 1.
Roberts, Craige. 1996. Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *In:* Yoon, J., & Kathol, A. (eds), *OSUWPL Volume 49: Papers in Semantics*.
Roberts, Craige. 2012. Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics*, **5**, 1–69.
Rothschild, Daniel. 2020. What It Takes to Believe. *Philosophical Studies*, **177**(5), 1345–1362.
Satta, Mark. 2018a. The Ambiguity Theory of "Knows". *Acta Analytica*, **33**(1), 69–83.
Satta, Mark. 2018b. A Linguistic Grounding for a Polysemy Theory of 'Knows'. *Philosophical Studies*, **175**(5), 1163–1182.
Satta, Mark. 2018c. Semantic Blindness and Error Theorizing for the Ambiguity Theory of 'Knows'. *Analysis*, **78**(2), 275–284.
Satta, Mark. 2020. Contextualism and the Ambiguity Theory of 'Knows'. *Episteme*, **17**(2), 209–229.
Sawyer, Sarah. 2014. Contrastive Self-knowledge. *Social Epistemology*, **28**(2), 139–152.
Schaffer, Jonathan. 2004. From Contextualism to Contrastivism. *Philosophical Studies*, **119**(1-2), 73–104.

Schaffer, Jonathan. 2005. Contrastive Knowledge. *Page 235 of:* Gendler, Tamar Szabo, & Hawthorne, John (eds), *Oxford Studies in Epistemology 1*. Oxford University Press.
Schaffer, Jonathan. 2007a. Closure, Contrast, and Answer. *Philosophical Studies*, **133**(2), 233–255.
Schaffer, Jonathan. 2007b. Knowing the Answer. *Philosophy and Phenomenological Research*, **75**(2), 383–403.
Schaffer, Jonathan. 2008. Knowledge in the Image of Assertion. *Philosophical Issues*, **18**(1), 1–19.
Schaffer, Jonathan. 2009. Knowing the Answer Redux: Replies to Brogaard and Kallestrup. *Philosophy and Phenomenological Research*, **78**(2), 477–500.
Schaffer, Jonathan. 2012. Contrastive Knowledge: Reply to Baumann. *Pages 411–24 of:* Tolksdorf, Stefan (ed), *The Concept of Knowledge*. Walter de Gruyter.
Schaffer, Jonathan, & Knobe, Joshua. 2012. Contrastive Knowledge Surveyed. *Noûs*, **46**(4), 675–708.
Schaffer, Jonathan, & Szabó, Zoltán Gendler. 2014. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies*, **168**(2), 491–543.
Schoenfield, Miriam. 2014. Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, **48**(2), 193–218.
Sosa, Ernest. 2015. *Judgment & Agency*. Oxford University Press UK.
Stanley, Jason. 2008. Knowledge and Certainty. *Philosophical Issues*, **18**(1), 35–57.
Steup, Matthias. 2005. Contextualism and Conceptual Disambiguation. *Acta Analytica*, **20**(1), 3–15.
Turri, John. 2010. Epistemic Invariantism and Speech Act Contextualism. *Philosophical Review*, **119**(1), 77–95.
Turri, John. 2011. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, **89**(1), 37–45.
Turri, John. 2015. Knowledge and the Norm of Assertion: A Simple Test. *Synthese*, **192**(2), 385–392.
Turri, John. 2016. *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Open Book Publishers.
Unger, Peter. 1968. An Analysis of Factual Knowledge. *Journal of Philosophy*, **65**(6), 157–170.
Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. Oxford: OUP.
White, Roger. 2005. Epistemic Permissiveness. *Philosophical Perspectives*, **19**(1), 445–459.
Williamson, Timothy. 1996. Knowing and Asserting. *Philosophical Review*, **105**(4), 489.
Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
Williamson, Timothy. 2011. Improbable Knowing. *In:* Dougherty, T. (ed), *Evidentialism and Its Discontents*. Oxford University Press.

Williamson, Timothy. 2013. Gettier Cases in Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(1), 1–14.

Williamson, Timothy. 2014. Very Improbable Knowing. *Erkenntnis*, **79**(5), 971–999.

Williamson, Timothy. forthcoming. Justifications, Excuses, and Sceptical Scenarios. *In:* Dorsch, Fabian, & Dutant, Julien (eds), *The New Evil Demon*. Oxford University Press.

Woozley, A. D. 1953. Knowing and Not Knowing. *Proceedings of the Aristotelian Society*, **53**, 151–172.

Zagzebski, Linda. 1994. The Inescapability of Gettier Problems. *Philosophical Quarterly*, **44**(174), 65–73.