Nonclassical Logic and Skepticism

Adam Marushak [*]

draft; please cite published version

Abstract

This paper introduces a novel strategy for responding to skeptical arguments based on the epistemic possibility of error or lack of certainty. I show that a nonclassical logic motivated by recent work on epistemic modals can be used to render such skeptical arguments invalid. That is, one can grant that knowledge is incompatible with the possibility of error and grant that error is possible, all while avoiding the skeptic's conclusion that we lack knowledge.

1 Introduction

Skeptical arguments often trade on the tension between claiming knowledge and allowing for the possibility of error. If one claims to know that ϕ, then it sounds contradictory to admit that ϕ might be false or that there remains uncertainty about whether ϕ. The skeptic presses us to concede that error is almost always possible, from which she concludes that we know very little.[1]

This paper introduces a novel strategy for responding to this type of skeptical argument. The existing replies in the literature generally come in one of the following three forms. The fallibilist argues that knowledge is compatible with the possibility of error.[2] The Moorean argues that we possess

---

[*]The author is grateful for helpful comments from Anil Gupta, James R. Shaw, two anonymous referees, and audiences at the Asian Epistemology Network and the University of Nevada at Las Vegas.

[1]Versions of this skeptical argument are discussed in Austin (1946), Moore (1959), Wittgenstein (1969), DeRose (1990, 1991), Lewis (1996), Engel (2004), Hawthorne (2004), Hill & Schechter (2007), Albritton (2011), and Reed (2013).

[2]See Dougherty & Rysiew (2009), Fantl & McGrath (2009*b*), and Reed (2013).

knowledge, so the possibility of error is an illusion.[3] The contextualist argues that consideration of skeptical possibilities changes the context so that the skeptic's conclusion is not at odds with ordinary claims to knowledge.[4] The strategy I introduce does something different: invariantists about 'knowledge' can grant that knowledge is incompatible with the possibility of error and grant that error is possible all while avoiding the skeptic's conclusion that we lack knowledge.

This strategy requires a nonclassical logic, for the skeptic's argument is valid if the consequence relation is classical: if knowledge is incompatible with the possiblity of error, then the possibility of error entails that one lacks knowledge (assuming there is no context shift mid-argument). I show, however, that there exists a well-motivated, nonclassical logic on which inferences of this type are not generally valid. The logic in question is drawn from recent work on epistemic modals. This paper shows that this logic has unexpected applications in epistemology that deserve further scrutiny.

2 Two Skeptical Arguments

Let's refine the skeptical arguments we alluded to above. Our focus will be on skeptical arguments based on the *epistemic* possibility of error—where epistemic possibility is the type of possibility characteristically expressed by modals like 'might', in its present tense, indicative form. For example, if one is wondering about the species of an animal in a zoo, one can say: "The animal might be a zebra." In this context, one can also express epistemic possibility by saying: "It is possible that the animal is a zebra." In other contexts, words like 'might' or 'possible' can be used to express metaphysical modality instead—although one generally needs to use the subjunctive form 'might have': "I might have chosen to be a lawyer instead of a philosopher." At the outset, we will avoid any theoretical characterization of epistemic modality in terms of knowledge (say). As we will see below, there are competing accounts of epistemic modality, some of which do not tie epistemic modality to knowledge or any epistemic state at all. For now, we are simply defining by ostension a flavor of natural language modality that figures in the skeptical arguments at issue.

[3] See Moore (1959).

[4] See DeRose (1990, 1995) and Lewis (1996).

Our first skeptical argument stems from the infelicity of so-called concessive knowledge attributions or CKAs—sentences of the form ⌜I know that ϕ, but it's possible/it might be not-ϕ⌝, where the modal in the second conjunct receives an epistemic reading.[5] For example, the following sounds contradictory or confused:

(1) # I know that the animal in the pen is a zebra, but it might not be a zebra.

We will schematize such sentences as follows: ⌜$K\phi$⌝ abbreviates ⌜I know that ϕ⌝, and $\Diamond$ represents natural language epistemic modals like 'might'. The logical form of (1) is thus given by (2):

(2) $K\phi \wedge \Diamond\neg\phi$

The infelicity of sentences with this form motivates the following principle, where '$\models$' represents the relation of consequence or entailment between sentences in natural language:

Contradiction: $(K\phi \wedge \Diamond\neg\phi) \models \bot$

Now, on a classical consequence relation, if a conjunction entails a contradiction, then each conjunct entails the negation of the other. Thus, if $\models$ obeys classical logic, then **Contradiction** entails a principle that Alex Worsnip (2015) calls the 'Knowledge-Possibility Link' or **KPL**:[6]

KPL: $\Diamond\neg\phi \models \neg K\phi$

With **KPL** in hand, the skeptic's next move is to claim that for many of the propositions we think we know, sober reflection forces us to concede that there remains at least some (epistemic) possibility that we are making a mistake. As Lewis (1996, 549) puts it:

> Let your paranoid fantasies rip—CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself—and soon you find that uneliminated possibilities of error are everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same.

[5]The label 'concessive knowledge attribution' is due to Rysiew (2001).

[6]Strictly speaking, Worsnip's **KPL** is a principle about truth at a context. I reframe the principle as a thesis about consequence.

Acknowledging that these possibilities of error remain open means that for virtually any sentence ϕ that one takes oneself to know, one must concede that *it might be* that not-ϕ. That is, the possibilities in question are characteristically epistemic, as Albritton (2011, 2) observes:

> The possibilities [the skeptical argument] alleges, from which it deduces rightly or wrongly that you don't know much, are such possibilities as that someone may be under the bed, or just might be under the bed. They are 'epistemic' possibilities, as we may say, 'remote,' perhaps, or 'bare' or 'faint,' certainly not necessarily 'live' or 'strong'—indeed, probably negligible, or typically neglected, at any rate, unless a lot turns on ruling them out—but not merely 'logical' or the like. And not what's currently called 'metaphysical,' either.

But once one concedes that it might be that not-ϕ (in the epistemic sense), the skeptic's conclusion follows, for by **KPL**, $\ulcorner\Diamond\neg\phi\urcorner$ entails that one does not know that ϕ. For example, if one concedes that the animal in the pen might not be a zebra (since it might be a cleverly painted mule), then **KPL** forces one to admit that one does not know that the animal is a zebra after all.

To generalize, we can think of the skeptic's argument as proceeding in two stages. In the first stage, the skeptic establishes the metalinguistic premise that sentences of the form $\ulcorner(K\phi \wedge \Diamond\neg\phi)\urcorner$ are contradictory, from which the skeptic infers the metalinguistic principle **KPL**. In the second stage, the skeptic establishes an object language claim of the form $\ulcorner\Diamond\neg\phi\urcorner$, from which the skeptic infers an object language claim of the form $\ulcorner\neg K\phi\urcorner$ via **KPL**. Thus, the skeptic's argument essentially rests on two premises—(P1) **Contradiction**; (P2) an object language epistemic possibility claim—and the lemma **KPL**, which is inferred from (P1) via the background assumption that $\models$ is classical.

Our second skeptical argument has a parallel structure. Adapting an argument from Unger (1975), Stanley (2008) observes that sentences like the following sound infelicitous:

(3) # I know that Bill came to the party, but it's not certain that he came to the party.

The second conjunct contains an expression of what I will call 'propositional certainty'—a proposition is described as being certain or not, as compared with describing a person as having certainty or not. We will abbreviate claims of propositional certainty as $\ulcorner C\phi \urcorner$, and so the logical form of (3) is given by (4):

(4) $K\phi \wedge \neg C\phi$

The infelicity of sentences with this form motivates an analogue to **Contradiction**:

Contradiction*: $(K\phi \wedge \neg C\phi) \models \bot$

If $\models$ obeys classical logic, then **Contradiction*** entails a principle that I will call the 'Knowledge-Certainty Link' or **KCL**:

KCL: $\neg C\phi \models \neg K\phi$

With **KCL** in hand, the skeptic invites us to reflect on whether what we take ourselves to know is indeed certain to be true. Lewis' 'paranoid fantasies' are once again in play. In fact, the considerations motivating the existence of uncertainty would seem to be exactly the same as those motivating the epistemic possibility of error. After all, it is highly plausible that claims of propositional certainty and epistemic possibility are duals: it might be that not-ϕ iff it is not certain that ϕ.[7] However, if we concede that what we think we know is not certain, the skeptic wins, for by **KCL**, $\ulcorner \neg C\phi \urcorner$ entails that I do not know that ϕ.

Each of these skeptical arguments is simple and powerful. The skeptic does not need to draw on any heavyweight, theoretical claims about the nature of knowledge, certainty, or epistemic possibility. The skeptic simply calls attention to the fact that (i) we ourselves seem to find it inconsistent to claim knowledge while allowing for the possibility of error or uncertainty, and (ii) we ourselves seem to concede that error or uncertainty always remains a possibility. Given a classical consequence relation, (i) and (ii) commit us to denying that we possess knowledge.[8]

[7] For endorsements of the duality thesis see DeRose (1998) and Littlejohn (2011).

[8] It is sometimes thought that the skeptic begs the question by assuming the epistemic possibility of error (see, for example, Hill & Schechter (2007, §5)). This is a mistake. The skeptic's point is that *we* are attracted to allowing for the epistemic possibility of error (witness the persistent attraction of fallibilism). The skeptic only aims to call attention to the inconsistency of our own position. It does not beg the question to point out that one's interlocutor has straightforwardly inconsistent beliefs.

3 A Comparison With Epistemic Modals

Suppose we apply the skeptic's reasoning to a similar piece of linguistic data. Sentences of the form $(\phi \wedge \Diamond\neg\phi)$—which Yalcin (2007) calls 'epistemic contradictions'—also sound bad:

(5) # It's raining, but it might not be raining.

The infelicity of epistemic contradictions motivates yet another **Contradiction** principle:

> **Contradiction****: $(\phi \wedge \Diamond\neg\phi) \models \bot$

If $\models$ obeys classical logic, then **Contradiction**** entails a principle that Worsnip (2015) calls the 'Truth-Possibility Link' or **TPL**:

> **TPL**: $\Diamond\neg\phi \models \neg\phi$

However, as Worsnip notes, something has gone very wrong here. **TPL** is absurd: epistemic modals are not factive operators. But where to locate the source of the problem?

One might first blame **Contradiction****: perhaps epistemic contradictions are not really contradictory (compare: the fallibilist argues that CKAs are not really contradictory either). After all, it is familiar that many sentences that sound infelicitous are not semantically contradictory, as Moore's paradox illustrates:

(6) # It's raining, but I don't know it.

(7) Suppose it's raining, but I don't know it.

The fact that it is coherent to entertain Moore-paradoxical sentences, as in (7), suggests that the infelicity of (6) is due to the pragmatics of assertion, not the sentence's being literally contradictory.

However, Yalcin (2007) observes that it is much harder to make this argument in the case of epistemic contradictions. He points out that epistemic contradictions sound bad even when embedded in environments where the sentences are not asserted:

(5) (repeated here) # It's raining, but it might not be raining.

(8) # Suppose it's raining, but it might not be raining.

Yalcin's observation is strong, prima facie evidence that epistemic contradictions are indeed semantically contradictory. His puzzle, then, is that each of the following claims is plausible but they cannot all be correct:

Classicality: $\models$ is classical

Nonfactivity: $\Diamond\neg\phi \not\models \neg\phi$

Contradiction**: $(\phi \wedge \Diamond\neg\phi) \models \bot$

Yalcin's (2007) preferred strategy is to reject **Classicality**. He develops an alternative theory of consequence he calls **Informational Consequence**, on which epistemic contradictions are contradictory but $\ulcorner\Diamond\neg\phi\urcorner$ does not entail $\ulcorner\neg\phi\urcorner$. We will explain how this theory works below. But for now, the key point is this: Yalcin's logic has exactly the properties we want in order to block the skeptical arguments discussed above. In both Yalcin's puzzle and our skeptical arguments, what we want is a logic wherein a conjunction is contradictory but it does not follow that each conjunct entails the negation of the other. A logic of this form allows Yalcin to grant that epistemic contradictions are genuinely contradictory while denying **TPL**. And a logic of this form would also allow us to grant that $\ulcorner(K\phi \wedge \Diamond\neg\phi)\urcorner$ and $\ulcorner(K\phi \wedge \neg C\phi)\urcorner$ are contradictory while denying **KPL** and **KCL**. The skeptic's arguments would thus be rendered invalid at the step from the **Contradiction** principles to the lemmas **KPL** and **KCL**. Alternatively, if one thinks of **Classicality** as a hidden premise in the skeptic's arguments, then our strategy is to expose and reject this hidden premise.

4 Informational Consequence

As a preliminary to our response to the skeptic, let us first explain how **Informational Consequence** resolves Yalcin's puzzle. In the next section we will show how this theory can also be used to block the skeptical arguments we discussed above.

Since Yalcin resolves his puzzle by rejecting **Classicality**, his task is to identify a nonclassical consequence relation on which $\ulcorner(\phi \wedge \Diamond\neg\phi)\urcorner$ is contradictory but $\ulcorner\Diamond\neg\phi\urcorner$ does not entail $\ulcorner\neg\phi\urcorner$. **Informational Consequence** has exactly these features, as we will explain below. The general idea behind this account of consequence is that sentences of a language place constraints on information states, where information states are understood as sets of worlds, and gaining information is understood as ruling out possibilities—that is, shrinking the set of worlds. On **Informational Consequence**, one sentence entails another iff every information state that satisfies the constraint placed by the first sentence satisifies the constraint placed by the second.

What does it mean for a sentence to place a constraint on an information state? Start with the case of a simple descriptive sentence like 'The animal is a zebra'. This sentence carries the information that the animal is a zebra, which is to say that the informational content of the sentence excludes all possibilities in which the animal is not a zebra. We can therefore think of the sentence as placing the following constraint on an information state: exclude all worlds in which the animal is not a zebra. An information state that satisfies this constraint will be one that contains only worlds at which the animal is a zebra.

But what is the constraint placed by an epistemically modalized sentence like 'The animal might be a zebra'? Here Yalcin departs from the orthodox view on which epistemic modalized sentences also carry information—information about how things stand with the speaker's knowledge (say).[9] That is, Yalcin denies that 'The animal might be a zebra' carries the information that the speaker's knowledge is consistent with the animal's being a zebra. He takes epistemically modalized sentences to place a different kind of constraint on information states: 'The animal might be a zebra' simply constrains an information state to be *compatible* with the animal's being a zebra. An information state might satisfy this constraint in a variety of different ways: perhaps the information state leaves open the possibility that the animal is a zebra or a mule, or perhaps the information state leaves open the possibility that the animal is a zebra or horse, and so on.

Yalcin formalizes this notion of an information state's 'satisfying a constraint' by what he calls *acceptance*. To understand his definition of acceptance we need to first review his definition of truth at a point of evaluation for epistemic modals:

> $[\![\Diamond\phi]\!]^{c,s,w} = 1$ iff $\exists w' \in s : [\![\phi]\!]^{c,s,w'} = 1$, where s is the so-called information state parameter, which ranges over sets of worlds, i.e. information

[9]For examples of the orthodox view, see DeRose (1991) and Stanley (2005).

states.

A few remarks on this notation. The double brackets denote the interpretation function, which outputs the semantic value of the sentence placed within the brackets. These semantic values can be evaluated at various parameters, such as a context c or world w. We use the '=1' notation in lieu of 'is true' in order to highlight the fact that a definition of truth at a point of evaluation does not necessarily track our intuitive notion of truth. In fact, on Yalcin's (2007, 2011) pragmatics, uses of epistemic modals do not express propositions that are properly assessed for truth or falsity at all. His definition of truth at a point of evaluation for epistemic modals is simply a device of giving a compositional semantics for epistemic modals, which can be used for generating semantic values for complex sentences embedding epistemic modals and also for defining a consequence relation for epistemic modals, as we show below.

On Yalcin's compositional semantics, epistemic modals are not sensitive to the world parameter w, as would be the case with an ordinary descriptive sentence like 'The animal is a zebra'. The truth of an epistemic modal is instead settled by the features of an information state s, which is given by a separate parameter. This compositional semantics is then put to work in Yalcin's definition of acceptance:

> s accepts a sentence-in-a-context ϕ_c iff $\forall w \in s : [\![\phi]\!]^{c,s,w} = 1$.

For an ordinary descriptive sentence that is not sensitive to the information state parameter, this definition says that an information state accepts that sentence iff all the worlds in the information state are worlds at which the sentence is true. For example, an information state accepts 'the animal is a zebra' iff all the worlds in the information state are worlds at which the animal is a zebra. But notice what happens in the case of an epistemically modalized sentence, which is sensitive to the information state parameter:

> s accepts $\Diamond\phi_c$ iff $\forall w \in s : [\![\Diamond\phi]\!]^{c,s,w} = 1$ iff $\forall w \in s : \exists w' \in s : [\![\phi]\!]^{c,s,w'} = 1$ iff $\exists w' \in s : [\![\phi]\!]^{c,s,w'} = 1$.

Since the truth of an epistemic modal at a point of evaluation turns only on the information state parameter s, not the world parameter w, the initial universal quantification over the worlds in the information state is redundant. The definition thus yields the desired result: an information state accepts an epistemically modalized sentence $\ulcorner\Diamond\phi\urcorner$ iff the information state is compatible with the proposition expressed by the modal prejacent ϕ. For example, an information state accepts 'The animal might be a zebra' iff the information state is compatible with the animal's being a zebra.

Finally, we give Yalcin's definition of **Informational Consequence**, which we now state in terms of preservation of acceptance:

> ϕ is an informational consequence of a set of sentences Γ ($\Gamma \models_I \phi$) iff for every context c and information state s, if s accepts ψ_c for every $\psi \in \Gamma$, then s accepts ϕ_c.

This definition formalizes the intuitive idea that one sentence entails another iff every information state that satisfies the constraint placed by the first also satisfies the constraint placed by the second.

To get a feel for how **Informational Consequence** works, consider the following inference:

$$\frac{\phi}{\neg\Diamond\neg\phi}$$

On an orthodox account of epistemic modals, this inference would be invalid. For example, an orthodox account might take epistemic modals to describe the knowledge of the speaker of the context and might follow Kaplan (1989) in taking consequence to be preservation of truth at a context. The above inference would then be invalid since there are contexts in which ϕ is true but in which the speaker of the context does not know that ϕ and so not-ϕ remains possible. **Informational Consequence** works differently. On this account, entailment is not based on any notion of truth at a context—in fact, Yalcin denies that epistemic modals have truth-conditions at a context at all.[10] Instead, entailment turns on relations between information states: every information state that accepts the premises accepts the conclusion. If an information state accepts ϕ, the information state has excluded all possibilities in which not-ϕ and thus it is not possible that not-ϕ. Hence, the above inference is valid.

[10] That is, Yalcin denies that there is such a thing as the information state of the context, s_c, that could be used to give a Kaplan-style definition of truth at a context for epistemically modalized sentences. The context parameter c functions in Yalcin's compositional semantics only to supply a proposition for context-sensitive sentences, where epistemic modals can then evaluate this proposition for compatibility with a given information state. Epistemic modals are not themselves context-sensitive.

We're now in a position to see how **Information Consequence** resolves Yalcin's puzzle: $(\phi \wedge \Diamond\neg\phi) \models_I \bot$ but $\Diamond\neg\phi \not\models_I \neg\phi$. $\ulcorner(\phi \wedge \Diamond\neg\phi)\urcorner$ is contradictory since no coherent information state accepts this sentence: accepting a conjunction requires accepting both conjuncts, but accepting the first conjunct requires that all the worlds in the information state are ϕ-worlds, while accepting the second conjunct requires that the information state is compatible with not-ϕ.[11]. Nevertheless, $\ulcorner\Diamond\neg\phi\urcorner$ does not entail $\ulcorner\neg\phi\urcorner$: there are information states that accept the first sentence but do not accept the second, such as an information state that contains both ϕ and not-ϕ worlds. This information state is compatible with not-ϕ and so accepts $\ulcorner\Diamond\neg\phi\urcorner$, but the information state does not accept $\ulcorner\neg\phi\urcorner$ since the information state also contains ϕ-worlds. Thus, **Informational Consequence** secures **Nonfactivity** and **Contradiction**** by giving up **Classicality**.

5 The Anti-Skeptical Strategy

We will now show how to use **Informational Consequence** to block the skeptical arguments we began with. Recall the parallel with Yalcin's puzzle and our skeptical arguments: in both cases, what we want is a logic wherein a conjunction is contradictory but it does not follow that each conjunct entails the negation of the other. We've just seen how **Informational Consequence** yields this result where $\ulcorner(\phi \wedge \Diamond\neg\phi)\urcorner$ is a contradiction but $\ulcorner\Diamond\neg\phi\urcorner$ does not entail $\ulcorner\neg\phi\urcorner$. Our task is to show how $\ulcorner(K\phi \wedge \Diamond\neg\phi)\urcorner$ and $\ulcorner(K\phi \wedge \neg C\phi)\urcorner$ can be contradictory, even while $\ulcorner\Diamond\neg\phi\urcorner$ does not entail $\ulcorner\neg K\phi\urcorner$, and $\ulcorner\neg C\phi\urcorner$ does not entail $\ulcorner\neg K\phi\urcorner$. A logic of this form would allow us to block the skeptic's inference from the **Contradiction** principles to **KPL** and **KCL**.

Consider the skeptic's first argument, which relies on the inference from CKAs being contradictory to **KPL**. Notice that on **Informational Consequence**, CKAs are contradictory for the same reason as epistemic contradictions: since knowledge is factive, any information state that accepts $\ulcorner K\phi\urcorner$ must contain only ϕ-worlds, in which case the information state cannot also accept $\ulcorner\Diamond\neg\phi\urcorner$, since the latter would require that the information state be compatible with $\ulcorner\neg\phi\urcorner$. But does it follow from this that $\ulcorner\Diamond\neg\phi\urcorner$ entails $\ulcorner\neg K\phi\urcorner$? The answer depends on the semantics of K.

[11] I speak of 'ϕ-worlds' for ease of exposition, but strictly speaking, points of evaluation are context/information-state/world triples.

Suppose $\ulcorner K\phi \urcorner$ were semantically equivalent to $\ulcorner \neg\Diamond\neg\phi \urcorner$, i.e. $\ulcorner \Box\phi \urcorner$:

$$[\![K\phi]\!]^{c,s,w} = 1 \text{ iff } \forall w' \in s : [\![\phi]\!]^{c,s,w'} = 1$$

On this semantics, **KPL** is indeed valid since $\Diamond\neg\phi \models_I \neg\Box\phi$ and $\ulcorner K\phi \urcorner$ is semantically equivalent to $\ulcorner \Box\phi \urcorner$. However, this semantics for K has an implausible consequence. **Informational Consequence** yields the result that $\Diamond\Box\phi \models_I \Box\phi \models_I \phi$.[12] Thus, since $\ulcorner K\phi \urcorner$ is semantically equivalent to $\ulcorner \Box\phi \urcorner$, $\Diamond K\phi \models_I K\phi \models_I \phi$. But this seems wrong: the mere epistemic possibility that I know that ϕ does not entail that I know that ϕ or that ϕ is true.[13]

Fortunately, there is an alternative account of the semantics of K that yields our desired result that **KPL** is invalid on **Informational Consequence**:

> $\Diamond\neg\phi \not\models_I \neg K\phi$ if there exists some context c and information state s such that $\exists w \in s : [\![\phi]\!]^{c,s,w} = 0$ and $\exists w \in s : [\![K\phi]\!]^{c,s,w} = 1$.

That is, **KPL** fails if it is possible for an information state to be undecided about whether ϕ and whether $\ulcorner K\phi \urcorner$. And this does seem possible: not every information state contains information as to exactly what the speaker of the context knows. An information state might be undecided as to whether an animal is a zebra and also whether I *know* that the animal is a zebra. Suppose we have an information state of this kind, where the information state is compatible with both $\ulcorner K\phi \urcorner$, ϕ, $\ulcorner \neg K\phi \urcorner$, and $\ulcorner \neg\phi \urcorner$. This information state is compatible with $\ulcorner \neg\phi \urcorner$ and so accepts $\ulcorner \Diamond\neg\phi \urcorner$, but the information state is compatible with $\ulcorner K\phi \urcorner$ and so does not accept $\ulcorner \neg K\phi \urcorner$. Hence, the information state provides a countermodel to **KPL**. The upshot: $\ulcorner (K\phi \wedge \Diamond\neg\phi) \urcorner$ is contradictory even though $\ulcorner \Diamond\neg\phi \urcorner$ does not entail $\ulcorner \neg K\phi \urcorner$. Thus, the skeptic's argument from **Contradiction** to **KPL** is invalid.

Let's reflect on the epistemological significance of this result. We've just supplied a logic wherein CKAs are contradictory but it does not follow from this fact that the epistemic possibility that not-ϕ entails that one lacks knowledge that ϕ. Interestingly, on our semantics, all that follows from the epistemic possibility that not-ϕ is that it is *epistemically possible* that one does

[12] It is known that on **Informational Consequence**, nested modals are semantically equivalent to the inner modal, so $\ulcorner \Diamond\Box\phi \urcorner$ is logically equivalent to $\ulcorner \Box\phi \urcorner$. $\ulcorner \Box\phi \urcorner$ entails ϕ since accepting $\ulcorner \Box\phi \urcorner$ involves redundant universal quantification over the information state parameter and is thus equivalent to accepting ϕ.

[13] See the last paragraph of this section and n. 16 for further discussion.

not know that ϕ. That is, since knowledge is factive, $\Diamond\neg\phi \models_I \Diamond\neg K\phi$: every information state that accepts $\ulcorner\Diamond\neg\phi\urcorner$ must contain at least one not-ϕ-world, which must also be a not-$K\phi$-world, and thus the information state must be compatible with $\ulcorner\neg K\phi\urcorner$ and hence accepts $\ulcorner\Diamond\neg K\phi\urcorner$. But this consequence does not give the skeptic her desired result, since $\ulcorner\Diamond\neg K\phi\urcorner$ does not entail $\ulcorner\neg K\phi\urcorner$ (a countermodel is again provided by an information state that is compatible with both $\ulcorner K\phi\urcorner$, ϕ, $\ulcorner\neg K\phi\urcorner$, and $\ulcorner\neg\phi\urcorner$).

Let's illustrate how this strategy works with an example. Suppose one grants to the skeptic that it is inconsistent to claim to know that an animal is a zebra while allowing that the animal might not be a zebra. Suppose one also grants to the skeptic that the animal might not be a zebra since the animal might be a painted mule. On our logic, it does not follow from this that one fails to know that the animal is a zebra. All that follows is that it is *epistemically possible* that one fails to know that the animal is a zebra. But it's being epistemically possible that one fails to know that the animal is a zebra does not entail that one *in fact* fails to know that the animal is a zebra. Hence, the skeptic's argument does not establish her intended conclusion that one does not know that the animal is a zebra.

Perhaps the skeptic has instead demonstrated that we fail to know that we know: if it might be that one does not know that ϕ, then one fails to know that one knows that ϕ. But notice that this argument relies on **KPL** as well: replacing ϕ in **KPL** with $\ulcorner K\phi\urcorner$, the skeptic infers that $\ulcorner\Diamond\neg K\phi\urcorner$ entails $\ulcorner\neg KK\phi\urcorner$. So by rejecting **KPL**, we also block the skeptic's argument against higher-order knowledge (a countermodel is provided by an information state that is compatible with $\ulcorner K\phi\urcorner$, $\ulcorner\neg K\phi\urcorner$, $\ulcorner KK\phi\urcorner$, and $\ulcorner\neg KK\phi\urcorner$).

Still, the skeptic does retain a small victory: since CKAs are contradictory, one cannot coherently assert that one knows the animal is a zebra while also asserting that the animal might not be a zebra. But of course, being unable to coherently assert a proposition is quite different from that proposition's being false. The skeptic has only succeeded in exploiting the dynamics of conversation to prevent us from properly *claiming* knowledge in certain contexts. ⌜I know that $\phi\urcorner$ may nevertheless be *true* even though it cannot be felicitously asserted in contexts where one has also asserted ⌜It might be that not-$\phi\urcorner$. Indeed, on anyone's theory, one cannot properly claim to know that ϕ while also asserting that it is possible that not-ϕ, since CKAs are at least pragmatically defective.[14] But surely the skeptic wishes to establish

[14]See Dougherty & Rysiew (2009) and Fantl & McGrath (2009*a*) for pragmatic expla-

more than the simple fact that CKAs are infelicitous to assert. Epistemic contradictions are also infelicitous to assert, but the infelicity of $\ulcorner \phi \wedge \Diamond\neg\phi \urcorner$ does not show that the possibility of error destroys truth, or that we can never properly assert that ϕ. Similarly, nothing in our reply to the skeptic entails that one can never properly assert that one knows that ϕ.

Now that we have the general structure of our anti-skeptical strategy in place, we can show how to extend this strategy to our second skeptical argument concerning the lack of certainty. Our strategy here is exactly the same: we will give a semantics for propositional certainty on which $\ulcorner (K\phi \wedge \neg C\phi) \urcorner$ is contradictory but $\ulcorner \neg C\phi \urcorner$ does not entail $\ulcorner \neg K\phi \urcorner$.

To do this, we will take propositional certainty to express epistemic necessity:

$$[\![C\phi]\!]^{c,s,w} = 1 \text{ iff } \forall w' \in s : [\![\phi]\!]^{c,s,w'} = 1.$$

That is, $\ulcorner C\phi \urcorner$ is semantically equivalent to $\ulcorner \neg\Diamond\neg\phi \urcorner$, and thus $\ulcorner \neg C\phi \urcorner$ is semantically equivalent to $\ulcorner \Diamond\neg\phi \urcorner$.[15] This semantics yields the result that $\ulcorner (K\phi \wedge \neg C\phi) \urcorner$ is contradictory for the same reason as $\ulcorner (K\phi \wedge \Diamond\neg\phi) \urcorner$: no coherent information state can accept both conjuncts—accepting $\ulcorner K\phi \urcorner$ requires the information state to contain only ϕ-worlds but accepting $\ulcorner \neg C\phi \urcorner$ requires the information state to contain at least one not-ϕ-world.

Does it follow from this that $\ulcorner \neg C\phi \urcorner$ entails $\ulcorner \neg K\phi \urcorner$? Since $\ulcorner \neg C\phi \urcorner$ is semantically equivalent to $\ulcorner \Diamond\neg\phi \urcorner$, the answer is the same as we saw above: it depends on the semantics of K. But again, we can block the entailment as long as it is possible for an information state to be undecided about both whether ϕ and whether $\ulcorner K\phi \urcorner$. An information state of this kind is compatible with $\ulcorner \neg\phi \urcorner$ and so accepts $\ulcorner \neg C\phi \urcorner$, but the information state is also compatible with $\ulcorner K\phi \urcorner$ and so does not accept $\ulcorner \neg K\phi \urcorner$.

In other words, a lack of certainty that ϕ simply leaves open the *possibility* that not-ϕ, which leaves open not-$K\phi$. But as we said above, epistemic modals are not factive: an information state's leaving open the possibility that one does not know that ϕ is different from an information state's *containing* the information that one does not know that ϕ. Our semantics does

nations of the infelicity of CKAs.

[15] A more realistic semantics for certainty would have to account for gradable certainty-talk (for example: $\ulcorner$It is fairly certain that $\phi\urcorner$). See Beddor (2020) for a semantics that handles gradable certainty-talk while remaining equivalent to a quantificational semantics in cases of non-graded certainty-talk.

yield the result that knowledge entails propositional certainty, but this stems only from the factivity of knowledge: since $\ulcorner K\phi \urcorner$ entails ϕ, $\ulcorner K\phi \urcorner$ excludes the possibility that not-ϕ and thus entails that it is certain that ϕ. But **Informational Consequence** does not validate contraposition: uncertainty about whether ϕ does not similarly contain the information that not-$K\phi$.

There is, however, one important difference between our response to the second skeptical argument concerning uncertainty and our response to the first concerning the possibility of error. As we noted above, $\Diamond\Box\phi \models_I \Box\phi \models_I \phi$. Since our semantics for C is identical to that of $\Box$, we are therefore committed to the result that $\Diamond C\phi \models_I C\phi \models_I \phi$. This result is problematic in the case of knowledge, but what about the case of certainty? Is the mere epistemic possibility that it is certain that ϕ enough to establish that it is certain that ϕ and hence establish that ϕ?

Here is one piece of data that suggests that our logic is correct. Notice that there seems to be a difference in felicity between the following:

(9) I think I know that ϕ, but it's possible that not-ϕ.

(10) ? I think it's certain that ϕ, but it's possible that not-ϕ.

There is something decidedly odd about (10): if one allows that it is possible that not-ϕ, one cannot at the same time hold out hope, as it were, that ϕ is certain to be true. In contrast, (9) sounds fine, as Worsnip (2015) notes. There seems to be nothing incoherent about allowing for the possibility that not-ϕ while holding out hope that one does in fact know that ϕ. For example, think of how the unconfident examinee might express her state of mind:

(11) I think I know the answer—A, but it might be B.

(12) ? I think it's certain that the answer is A, but it might be B.

The data here is subtle, but if these intuitions are correct, they support the result that $\Diamond C\phi \models_I C\phi \models_I \phi$. For suppose that ⌜I think that ϕ⌝ commits the speaker to $\ulcorner \Diamond\phi \urcorner$. It would then follow that ⌜I think it's certain that ϕ, but it's possible that not-ϕ⌝ conveys $\ulcorner \Diamond C\phi \wedge \Diamond\neg\phi \urcorner$, which by the aforementioned result entails the epistemic contradiction $\ulcorner \phi \wedge \Diamond\neg\phi \urcorner$, which is contradictory on **Informational Consequence**. In contrast, $\ulcorner \Diamond K\phi \wedge \Diamond\neg\phi \urcorner$ is consistent, assuming again that there are information states that are compatible with both $\ulcorner K\phi \urcorner$, ϕ, $\ulcorner \neg K\phi \urcorner$, and $\ulcorner \neg\phi \urcorner$.[16]

[16]Note that the intuitive acceptability of $\ulcorner \Diamond K\phi \wedge \Diamond\neg\phi \urcorner$ is additional evidence that

6 Conclusion

We've laid out proof of concept for a novel anti-skeptical strategy based on nonclassical logic. I close by reviewing some of the outstanding questions surrounding this approach.

First, our anti-skeptical strategy is directed at two skeptical arguments that involve an inference from a conjunction's being contradictory to one conjunct's entailing the negation of the other. It is an open question whether all skeptical arguments appealing to the epistemic possibility of error have this form. In particular, one wonders whether the skeptic could move directly to motivating principles like **KPL** or **KCL** without first inferring them from **Contradiction** principles. However, even if the skeptic's arguments are reformulated in this way, our logic will still be of interest, since it shows how it is possible to give semantics for knowledge, certainty, and epistemic possibility on which **KPL** and **KCL** are both false, and so the skeptic's arguments remain unsound.

Second, **Informational Consequence** exploits distinctive features of the semantics of epistemic modals. As it turns out, the theory respects classical logic when it comes to inferences that do not involve epistemic modals. Thus, it is not clear that **Informational Consequence** can be extended to block skeptical arguments that do not contain epistemic modals. Examples of such arguments include skeptical arguments based on the metaphysical possibility of skeptical hypotheses, not their epistemic possibility. Many other skeptical arguments do not reference any type of possibility at all, such as familiar closure-based arguments for skepticism. Nevertheless, it is not clear that we should expect a unified response to every type of skeptical argument. It is enough if our anti-skeptical strategy suffices to block at least one important type of skeptical argument based on epistemic modals.

Third, **Informational Consequence** is controversial, and we have not attempted to assess the arguments for and against this way of understanding the natural language consequence relation.[17] Our response to the skeptic thus remains tenative, and it is open to the skeptic to resist our response by arguing that the consequence relation is classical. However, for the skeptic to make this move would already be a significant concession. After all,

$\ulcorner \Diamond K\phi \urcorner$ does not entail ϕ, since if it did, then $\ulcorner \Diamond K\phi \wedge \Diamond\neg\phi \urcorner$ would entail the intuitively unacceptable $\ulcorner \phi \wedge \Diamond\neg\phi \urcorner$.

[17]For critical assessment of **Informational Consequence** see Schulz (2010), Bledin (2014), and Santorio (2022).

the skeptic's arguments were supposed to be based only on intuitive considerations about knowledge and possibility—not some theoretical account of the natural language consequence relation. Moreover, Yalcin and others have already given independent motivation for thinking that **Informational Consequence** is the correct account of consequence in natural language.[18] There may also be ways of developing analogous responses to the skeptic within the framework of other, independently-motivated nonclassical consequence relations. For example, we noted that the skeptic's arguments rest on the assumption that if a conjunction is contradictory, then each conjunct entails the negation of the other, and we saw that **Informational Consequence** falsifies this assumption. But other nonclassical accounts of consequence also falsify this assumption.[19] So there may well exist several routes to a nonclassical response to the skeptic, not all of which rely on **Informational Consequence** in particular.

Finally, we have not examined how our nonclassical strategy compares to the anti-skeptical strategies of the fallibilist, the Moorean, and the contextualist. What are the distinctive costs and benefits of the various approaches? I leave this question and the others as topics for future research.

References

Albritton, R. (2011), 'On a form of skeptical argument from possibility', *Philosophical Issues* **21**(1), 1–24.

Austin, J. (1946), 'Other minds', *Aristotelian Society Supplementary Volume* **20**, 148–87.

Beddor, B. (2020), 'New work for certainty', *Philosophers' Imprint* **20**(8).

Bledin, J. (2014), 'Logic informed', *Mind* **123**(490), 277–316.

Bledin, J. & Lando, T. (2018), 'Closure and epistemic modals', *Philosophy and Phenomenological Research* **97**(1), 3–22.

DeRose, K. (1990), Knowledge, Epistemic Possibility, Scepticism, PhD thesis, UCLA.

[18] See Yalcin (2007, section 4), Bledin (2014), and Bledin & Lando (2018).

[19] See, for example, Veltman's (1996) definitions of consequence in dynamic semantics and Mandelkern's (2019) bounded theory.

DeRose, K. (1991), 'Epistemic Possibility', *Philosophical Review* **100**, 581–605.

DeRose, K. (1995), 'Solving the skeptical problem', *Philosophical Review* **104**(1), 1–52.

DeRose, K. (1998), 'Simple 'might's, indicative possibilities and the open future', *Philosophical Quarterly* **48**(190), 67–82.

Dougherty, T. & Rysiew, P. (2009), 'Fallibilism, epistemic possibility, and concessive knowledge attributions', *Philosophy and Phenomenological Research* **78**(1), 123–132.

Egan, A. & Weatherson, B., eds (2011), *Epistemic Modality*, Oxford University Press, Oxford.

Engel, M. (2004), 'What's wrong with contextualism, and a noncontextualist resolution of the skeptical paradox', *Erkenntnis* **61**(2-3), 203–231.

Fantl, J. & McGrath, M. (2009*a*), 'Advice for fallibilists: Put knowledge to work', *Philosophical Studies* **142**(1), 55–66.

Fantl, J. & McGrath, M. (2009*b*), *Knowledge in an Uncertain World*, Oxford University Press.

Hawthorne, J. (2004), *Knowledge and Lotteries*, Oxford University Press.

Hill, C. S. & Schechter, J. (2007), 'Hawthorne?s lottery puzzle and the nature of belief', *Philosophical Issues* **17**(1), 120–122.

Kaplan, D. (1989), Demonstratives: An essay on the semantics, logic, metaphysics and epistemology of demonstratives and other indexicals, *in* J. Almog, J. Perry & H. Wettstein, eds, 'Themes From Kaplan', Oxford University Press, pp. 481–563.

Lewis, D. (1996), 'Elusive knowledge', *Australasian Journal of Philosophy* **74**(4), 549–567.

Littlejohn, C. (2011), 'Concessive knowledge attributions and fallibilism', *Philosophy and Phenomenological Research* **83**(3), 603–619.

Mandelkern, M. (2019), 'Bounded modality', *Philosophical Review* **128**(1), 1–61.

Moore, G. E. (1959), *Philosophical Papers*, New York: Routledge.

Reed, B. (2013), 'Fallibilism, epistemic possibility, and epistemic agency', *Philosophical Issues* **23**(1), 40–69.

Rysiew, P. (2001), 'The context-sensitivity of knowledge attributions', *Noûs* **35**(4), 477–514.

Santorio, P. (2022), 'Trivializing informational consequence', *Philosophy and Phenomenological Research* **104**(2), 297–320.

Schulz, M. (2010), 'Epistemic modals and informational consequence', *Synthese* **174**(3), 385–395.

Stanley, J. (2005), 'Fallibilism and Concessive Knowledge Attributions', *Analysis* **65**(2), 126–31.

Stanley, J. (2008), 'Knowledge and certainty', *Philosophical Issues* **18**(1), 35–57.

Unger, P. K. (1975), *Ignorance: A Case for Scepticism*, Oxford University Press.

Veltman, F. (1996), 'Defaults in update semantics', *Journal of Philosophical Logic* **25**(3), 221–261.

Wittgenstein, L. (1969), *On Certainty (Ed. Anscombe and von Wright)*, New York and London: Harper Torchbooks.

Worsnip, A. (2015), 'Possibly false knowledge', *Journal of Philosophy* **112**(5), 225–246.

Yalcin, S. (2007), 'Epistemic Modals', *Mind* **116**(464), 983–1026.

Yalcin, S. (2011), Nonfactualism about Epistemic Modality, *in* '(**?**)', pp. 295–332.