Closure and Epistemic Modals

Justin Bledin and Tamar Lando

July 16, 2015

Abstract: According to a popular closure principle for epistemic justification, if one is justified in believing that each premise in Φ is true and one comes to believe that ψ is true on the basis of competently deducing ψ from Φ—while retaining justified beliefs in the premises—then one is justified in believing that ψ is true. This principle is *prima facie* compelling; it seems to capture the sense in which competent deduction is an epistemically secure means to extend belief. Nevertheless, we argue that when closure is applied to our ordinary everyday concept of belief, even the single-premise version of this principle fails. Our counterexamples involve the epistemic possibility operator $\Diamond$. Though one can competently deduce $\neg\Diamond\varphi$ from $\neg\varphi$ in deliberation, there are cases in which one can justifiably believe that $\neg\varphi$ is true but would be unjustified in believing that $\neg\Diamond\varphi$ is true.

1 Against Closure

It is not the case that φ is true; therefore, it is not the case that φ might be true. When the possibility modal in this schema is given an epistemic interpretation, linguists and philosophers of language like Veltman [1996] and Yalcin [2007] have offered semantic accounts that validate arguments of this form. But if you can competently make such arguments within deliberation, this threatens closure for epistemic justification.

Here is why. According to a naïve closure principle for justification (refined below), you can justifiably believe what you competently infer from your justified beliefs. For some φ, though, we will argue that there are deliberative contexts in which you are justified in believing that it is not the case that φ is true but would be unjustified in believing that it is not the case that φ might be true—at least there are such contexts

when *belief* is understood in the ordinary everyday sense. So closure for epistemic justification is in trouble.

Consider this example. It is the lead-up to the 1980 U. S. presidential election and opinion polls project that Ronald Reagan will win handily, with Jimmy Carter coming in second place and John Anderson coming a distant third. Suppose that you come to believe, on the basis of these polls, that the following sentence is true:

(1) Carter will not win the election.

Arguably, this belief is justified. However, given that your evidence does not conclusively rule out the possibility of a Carter win, you would arguably be unjustified in, or have insufficient grounds for, believing that this modalized sentence is true:

(2) It is not the case that Carter might win the election.

Many similar examples can be constructed involving the prediction of future events about which you have no certain knowledge.[1]

Another class of problematic arguments with the same form concerns past events with low objective chance. Suppose that a fair lottery with one thousand tickets has been held but you are unaware of the result. In a deliberative context where you are considering, say, whether ticket 10 won, you are arguably justified in believing that this sentence is true:

(3) Ticket 10 did not win.[2]

---

[1] As some readers will have recognized, this example is a spinoff of McGee's [1985] famous 'counterexample' to *modus ponens* for the indicative conditional:

(P1) If a Republican wins the election, then if it is not Reagan who wins it will be Anderson.

(P2) A Republican will win the election.

(C) If it is not Reagan who wins, it will be Anderson.

McGee and many commentators on his paper agree that one can justifiably believe, on the basis of the opinion polls, that (P2) is true; presumably, they will also agree that one can justifiably believe that Carter will not win the election. Moreover, McGee and commentators agree that one would be unjustified in believing that (C) is true; presumably, they will also agree that one would be unjustified in believing that it is not the case that Carter might win.

[2] Admittedly, it is controversial whether you are justified in believing this premise on the basis of your purely statistical evidence. Stalnaker [1984], among many others, argues that you are justified only in having a very high credence, but not a full belief, that ticket 10 did not win. We will have more to say about this in what follows. For now, let us just point out that it is enough for the purpose of our argument that you are justified in believing the premise but would be unjustified in believing the conclusion in either of our examples.

But you would be unjustified in believing that this sentence is true:

(4) It is not the case that ticket 10 might have won.

After all, ticket 10 is just as likely to have won as any other.

Together with preface cases, such lottery cases have often been taken to undermine a *multi-premise* closure under conjunction principle for epistemic justification. It is commonly held that you can justifiably believe that ticket i did not win for each $i \leq 1000$ but you would be unjustified in believing the conjunction that tickets 1 through 1000 did not win.[3] It is also commonly held that a historian can justifiably believe each of the many claims made in one of her books but she would be unjustified in believing that their conjunction is true (or that the book is error-free).[4] Since justified belief is compatible with some risk of inaccuracy, the story goes, this risk can aggregate over a multi-premise deduction and undermine justified belief in its conclusion. But our two examples, by contrast, are proposed counterexamples to *single-premise* closure. The loss of justification in our examples does not result from the aggregation of risk across multiple premises.[5]

Note that closure principles for epistemic justification must be sharply distinguished from related principles for knowledge. According to a naïve closure principle for knowledge, you know what you competently deduce from known premises. One can hold that knowledge is closed under competent deduction without holding that justification is closed under competent deduction, and vice versa.[6] Indeed, our two examples are *not* counterexamples to a single-premise closure principle for knowledge. If you *know* that Carter will not win the election, then you presumably also know that it is not the case that he might win. Similarly, if you *know* that ticket 10 did not win the lottery, then you presumably also know that it is not the case that this ticket might have won.

[3]In fact, putting your cognitive limitations to the side, you arguably ought to believe the negation of this conjunction. The lottery paradox was introduced by Kyburg [1961] to demonstrate, *inter alia*, that rational acceptance is not closed under conjunction. See Wheeler [2007] for a comprehensive review of the vast literature on this paradox.

[4]The preface paradox was introduced by Makinson [1965]. See Christensen [2004] for a good extensive discussion of this paradox.

[5]Schechter [2013] also challenges single-premise closure for epistemic justification. But unlike our attack, his "Long Sequence Argument" relies on the phenomenon of *rational self-doubt*: if you start with a single justifiably believed premise and then competently perform a very long chain of deductions, you would be unjustified in believing its ultimate conclusion given that we are all prone to error in our reasoning.

[6]Lewis [1996], for instance, defends a closure principle for knowledge but claims that one is justified in believing of each ticket in a large lottery that it will lose, so he presumably rejects a general closure principle for justification.

2 A Threefold Tension

The problem for closure can be put more neutrally in terms of a tension between three principles. The first is single-premise closure for epistemic justification, here formulated with a bit more care:

Single-Premise Closure: For any sentences φ and ψ, if one is justified in believing that φ is true and one comes to believe that ψ is true on the basis of competently deducing ψ from φ—while justifiably retaining one's belief that φ is true—then one is justified in believing that ψ is true.

Which concept of 'justification' enters into this principle? We agree with Schechter [2013] that the notion of justification that supports closure has to do with *epistemic responsibility*:

> The central intuition supporting closure is that deduction is a *responsible* belief-forming method. Thinkers are epistemically responsible in believing what they deductively infer from epistemically responsible beliefs. (p. 433)

According to Single-Premise Closure, coming to believe that φ is true through some epistemically responsible process and then competently inferring ψ from φ is itself a responsible method for forming the belief that ψ is true.

The second principle concerns good inference:

Łukasiewicz's Principle: For each φ, one can competently deduce $\neg\Diamond\varphi$ from $\neg\varphi$ in any deliberative context.

We name this principle after Jan Łukasiewicz because, as Yalcin [2007] reports, Łukasiewicz [1930] seems to endorse it—at least in hypothetical contexts where you are supposing that $\neg\varphi$ is true.[7]

The third principle is this:

Justification with Risk: For some φ, there are deliberative contexts in which one can justifiably believe that $\neg\varphi$ is true but one would be unjustified in believing that $\neg\Diamond\varphi$ is true.

These principles clearly conflict. Consider a sentence φ that witnesses Justification with Risk and suppose that you justifiably believe that $\neg\varphi$ is true in a deliberative context where you would be unjustified in

[7] Our terminology differs from Yalcin's in that his 'Łukasiewicz's Principle' states that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are *valid* whereas our version states only that one can competently make such arguments in both categorical and hypothetical deliberative contexts. More on this in §3.

believing that $\neg\Diamond\varphi$ is true. Now suppose that you infer $\neg\Diamond\varphi$ from $\neg\varphi$ while retaining your justified belief. By Łukasiewicz's Principle, you have performed a competent deduction. By Single-Premise Closure, you are justified in believing that $\neg\Diamond\varphi$ is true. So something has to give. In the following two sections, we defend Łukasiewicz's Principle and Justification with Risk. This leads us to reject Single-Premise Closure.

3 Łukasiewicz's Principle

We suspect that Łukasiewicz's Principle will meet with opposition. After all, it is widely accepted that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are generally invalid. On the traditional *contextualist* semantics for epistemic modals endorsed in various forms by Hacking [1967], Teller [1972], Kratzer [1981], DeRose [1991], Dowell [2011], von Fintel and Gillies [2011], Yanovich [2014], and many others, the truth of an epistemic modal sentence turns on what some contextually relevant individual or group of individuals knows or can come to know through certain channels of investigation—hence the standard moniker 'epistemic modal'. Specifically, $\Diamond\varphi$ is true in a context just in case the salient epistemic state supplied by this context leaves open the possibility that φ is true. Also, as ordinarily conceived, an argument is valid just in case there is no context in which each of the premises are true but the conclusion is false. Putting these two pieces together: for many φ, the argument from $\neg\varphi$ to $\neg\Diamond\varphi$ is invalid since there are contexts in which φ is false but the knowledge or potential knowledge state relevant to the evaluation of epistemic modal sentences leaves open the possibility of its truth.

Now, by itself, this does not tell against Łukasiewicz's Principle. This principle states not that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are valid but only that each argument of this form is a good argument that a reasoner can competently make in any categorical deliberative context where she knows its premise and in any hypothetical context where she is only supposing it. As we understand these two notions, *good argument* and *validity* are conceptually distinct. The former is a largely pre-theoretic evaluative concept applicable in the third-person standpoint of appraisal, the perspective in which even those without formal logic training will evaluate arguments as good or bad and the activity of making them as competent or incompetent. The latter is a theoretical concept concerned with truth preservation or some other descriptive feature of arguments; it alone says nothing about what we can and cannot infer in the course of our deliberations.[8] This being said, many philosophers still assume

[8]See Bledin [2015] for related comparison of *deductively good* arguments, *logically*

that these different notions line up closely, that an argument is valid if and only if one can competently make it by virtue of the meaning of its constituent expressions. So the invalidity of our target arguments on the traditional picture seems to provide good grounds for rejecting Łukasiewicz's Principle.

To defend our principle, we could push the line that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are in fact valid. We would certainly not be alone in doing so. Though he does not explicitly discuss arguments of this form, Veltman [1996] presents a dynamic "update semantics" for epistemic modals and defines various consequence relations over it each of which validates these arguments (his dynamic semantics has subsequently been taken up by Beaver [2001], Willer [2013], Starr [2014], and many others).[9] Yalcin [2007] also presents an "informational consequence" relation defined over

---

valid arguments, and arguments *that are normative for thought*.

[9]For those readers interested in the formal details, here is a quick run-through of Veltman's semantics. The basic idea behind dynamic semantics is that the meaning of a sentence is not its truth conditions but rather a program or instruction for updating information (its *context change potential*). Motivated by Stalnaker's [1978] classic conception of assertions as proposals to update *common ground*, Veltman compositionally pairs sentences in a language with functions from information states (modeled using sets of possible worlds) to information states. A simple language $\mathcal{L}$ with only sentence letters A, B, ..., negation $\neg$, and the epistemic possibility modal $\Diamond$ is interpreted as follows:

Def. A *model* $\mathcal{M} = \langle \mathcal{W}, \mathcal{V} \rangle$ for $\mathcal{L}$ consists of a nonempty set of worlds $\mathcal{W}$ and an interpretation function $\mathcal{V}$ that maps each sentence letter and world $w \in \mathcal{W}$ to 1 or 0.

Def. The *update function* $[\![\ ]\!]_{\mathcal{M}}$ for $\mathcal{L}$ maps each sentence to a function that in turn maps each information state $i \in 2^{\mathcal{W}}$ to an information state (the clause for A is representative):

$$\begin{array}{rcl} i[\![A]\!]_{\mathcal{M}} & = & \{w \in i : \mathcal{V}(A, w) = 1\} \\ i[\![\neg\varphi]\!]_{\mathcal{M}} & = & i \setminus i[\![\varphi]\!]_{\mathcal{M}} \\ i[\![\Diamond\varphi]\!]_{\mathcal{M}} & = & \{w \in i : i[\![\varphi]\!]_{\mathcal{M}} \neq \emptyset\} \end{array}$$

Updating i with A amounts to removing every A-world from i. Updating i with $\neg\varphi$ amounts to first updating i with φ and then removing worlds in the posterior state $i[\![\varphi]\!]_{\mathcal{M}}$ from the prior state i. Lastly, $\Diamond$ tests whether updating i with φ fails to return $\emptyset$. If this test passes, then the posterior state $i[\![\Diamond\varphi]\!]_{\mathcal{M}}$ is the input state i itself; otherwise, the output is $\emptyset$.

Veltman defines a few different consequence relations over this dynamic semantics. Here is one of them:

Def. $i \rhd \varphi$ (read: i *incorporates* or *supports* φ) iff $i[\![\varphi]\!]_{\mathcal{M}} = i$.

Def. $\{\varphi_1, ..., \varphi_n\} \models \psi$ iff there is no $\mathcal{M}$ such that for some $i \in 2^{\mathcal{W}}$, $i \rhd \varphi_1, ..., i \rhd \varphi_n$ but $i \not\rhd \psi$.

That is, an argument is valid iff every information state in a model that incorporates its premises also incorporates its conclusion. It is easy to verify that $i \rhd \neg\varphi$ implies $i \rhd \neg\Diamond\varphi$, so $\{\neg\varphi\} \models \neg\Diamond\varphi$.

a static semantics that validates arguments from $\neg\varphi$ to $\neg\Diamond\varphi$.[10] But while we find these accounts compelling, we rest content to defend the weaker position here that our target arguments are good ones to make in deliberation. We leave open the possibility that these arguments are nevertheless invalid, and so we allow for the notions of good argument and validity to potentially come apart.

Why think that the Łukasiewicz arguments are good arguments? Well, as discussed in Bledin [2013], [2014], [2015], Yalcin's informational consequence relation (and Veltman's dynamic analogue of this relation) motivates an informational conception of deductive inquiry that makes good sense of our intuitions regarding competent deduction in modal and conditional languages. Even if one rejects the informational account of consequence, one can still accept the informational account of deduction. The basic idea underlying this account is that deductive inquiry is an information-driven enterprise in which a reasoner investigates what is so according to some salient information—typically the content of her beliefs—that *incorporates* the premises of some argument. For example, suppose that a market report tells us the following:

(5) The price of soybeans might rise.

(6) If the price of soybeans rises, the supply of corn will increase.

(7) The demand for sorghum will not increase.

What else does this report incorporate?

To make this a bit more precise, think of a body of information as that which rules out various possible ways the world might be, or might have been, while leaving others open. Information will then incorporate that such and such just in case it has a particular possible-worlds based structure. For example: the market report incorporates that (5) is true since it leaves open the possibility that the price of soybeans will rise, the report incorporates that (6) is true since it rules out the possibility that the price of soybeans will rise but the supply of corn will not, and it incorporates that (7) is true since it rules out the possibility that the demand for sorghum will increase.[11] Reasoning with this information,

[10] Yalcin's informational consequence relation is effectively a staticized version of Veltman's relation $\models$.

[11] The technical relation $\rhd$ defined in n. 9 is intended to explicate this informal notion of incorporation. Let A, B, and C abbreviate 'The price of soybeans will rise', 'The supply of corn will increase', and 'The demand for sorghum will increase' respectively, and let us supplement Veltman's recursive semantics with the following clause for the indicative conditional $\Rightarrow$ (Gillies [2010]):

$$i[\![\varphi \Rightarrow \psi]\!]_{\mathcal{M}} \quad = \quad \{w \in i : i[\![\varphi]\!]_{\mathcal{M}}[\![\psi]\!]_{\mathcal{M}} = i[\![\varphi]\!]_{\mathcal{M}}\}$$

the inferences that one can competently draw on its basis are precisely those that preserve incorporation. Since the market report incorporates that (5) and (6) are both true, it leaves open the possibility that the supply of corn will increase, so it incorporates the following:

(8) The supply of corn might increase.[12]

Since the report rules out the possibility that the demand for sorghum will increase, it also incorporates the following:

(9) It is not the case that the demand for sorghum might increase.

More generally, information incorporating that $\neg\varphi$ is true is therefore also information incorporating that $\neg\Diamond\varphi$ is true, so one can competently infer the latter from the former in deliberation. This is, of course, just what Łukasiewicz's Principle tells us.

Again, philosophers who cling to the traditional contextualist view of epistemic modals and the truth preservation view of validity discussed at the beginning of this section will deny that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are valid.[13] But so long as they accept the informational account of deduction just sketched, they should accept Łukasiewicz's Principle. The informational account is supported by two sorts of data. On the one hand, many intuitively good arguments that standard truth-based accounts tell us are bad count as good arguments on the informational view. For example: *modus ponens* for the indicative conditional and related arguments like the one from (5) and (6) to (8) presented above.[14] On the other hand, the informational view also predicts and explains the badness we find with certain applications of *modus tollens* (Yalcin [2012a]) and with certain episodes of reasoning involving *reductio* and proof by cases in languages with informational modals and the indicative conditional (Kolodny and MacFarlane [2010]). See Bledin [2014], [2015] for more extensive discussion.

One final point. Our distinction between validity and good argument might bring to mind Stalnaker's [1975] well-known distinction between

---

The agricultural report can then be modeled with an information state $i \in 2^{\mathcal{W}}$ where $i \rhd \Diamond A$, $i \rhd A \Rightarrow B$, and $i \rhd \neg C$. Assuming that i is nonempty, it includes at least some A-worlds, includes only A-worlds that are also B-worlds, and excludes C-worlds.

[12]Using the semantics in n. 9 and n. 11, it is easy to verify that $i \rhd \Diamond A$ and $i \rhd A \Rightarrow B$ together imply $i \rhd \Diamond B$, so $\{\Diamond A, A \Rightarrow B\} \models \Diamond B$.

[13]Yalcin [2011] himself argues that epistemic modalized sentences like $\neg\Diamond\varphi$ are not even, properly speaking, true or false at contexts of use. On his *non-factualist* position, Łukasiewicz arguments fail to preserve truth in a particularly strong sense.

[14]Bledin [2015] defends the view that *modus ponens* for the indicative is a good form of argumentation against multiple attacks.

the semantic concept of *entailment*—which he understands in terms of truth preservation—and the pragmatic concept of *reasonable inference*. According to Stalnaker, for instance, the *or-to-if* inference from 'Either the butler or the gardener did it' to 'If the butler didn't do it then the gardener did' is invalid since the former sentence does not semantically entail the latter, but this inference is still reasonable since if one can appropriately assert the disjunctive premise in a context of use, then the background *common ground* information in this context will incorporate the conditional conclusion. However, while the concepts of reasonable inference and good argument are closely related, they are not the same. While every good argument is reasonable, there are reasonable inferences that are not good.[15] More significantly, our notion of good argument is not tied to the pragmatics of assertion. One can talk of good arguments in the the private contexts of our internal deliberations without having to posit some peculiar 'pragmatics of thought'.

4 Justification with Risk

Moving on to Justification with Risk, we will now argue that there are cases in which you are justified in believing that $\neg\varphi$ is true but you would be unjustified in believing that $\neg\Diamond\varphi$ is true—at least such cases exist when working with our basic ordinary concept of belief. Indeed, we will argue that our two examples from §1 fit this mold. In the election case, you are justified in believing that Carter will not win the election but you would be unjustified in believing that it is not the case that Carter might win. In the lottery case, you are justified in believing that ticket 10 did not win the lottery but you would be unjustified in believing that it is not the case that ticket 10 might have won.

Our argument in this section has two parts. First, we will argue that there are cases in which one can justifiably believe—in the ordinary sense of 'believe'—that $\neg\varphi$ is true even though one's available evidence leaves open the possibility that φ is true and one knows that it does. Next, we will argue that in at least some of these cases, one would be unjustified in believing that $\neg\Diamond\varphi$ is true. Putting these two pieces together, our defense of Justification with Risk will be complete.

Since we are dealing with our basic concept of belief, we can look to natural language belief reports to better understand it. Our first bit of linguistic data comes from Hawthorne, Rothschild, and Spectre [*ms.*]:

(10) # Tim thinks it's raining, but he doesn't believe that it is.

[15] See Bledin [2014], n. 29 for an example.

(11) # Tim is of the opinion that it will rain, but he doesn't go so far as to believe that it will.

Both of these attitude reports are contradictory-sounding. While there might be some creative pragmatic explanation of their oddity, we cannot see how this would go.[16] So we agree with Hawthorne, Rothschild, and Spectre that this data shows that thinking or being of the opinion that it is raining is a sufficient condition for believing that it is. The following sentences also sound terrible:

(12) # Tim believes that it's raining, but he doesn't think it is.

(13) # Tim believes that it will rain, but he is not of the opinion that it will.

This presumably shows that thinking or being of the opinion that it is raining is also a necessary condition for believing that it is. The more general conclusion for at least non-modal φ is this: believing that φ is true coincides with such epistemic states as thinking and being of the opinion that φ is true. One cannot be in any one of these states without being in the others.[17]

Partly on the basis of the above data, Hawthorne, Rothschild, and Spectre conclude that "belief is weak". That is, they conclude that the evidential warrant for believing that φ is true is as low as the evidential warrant for thinking or being of the opinion that φ is true. At least when φ is non-modal, we agree with this. So even without saying more about what it is exactly to believe in the ordinary sense, and about the exact conditions under which one can responsibly acquire beliefs, we maintain that there are cases in which one can justifiably believe that $\neg\varphi$ is true despite knowing that one's total available evidence leaves open the possibility that φ is true. After all, if one's recognizably non-conclusive evidence still overwhelmingly supports that $\neg\varphi$ is true, then surely one can justifiably think or be of the opinion that $\neg\varphi$ is true. In the election case, you know that your evidence does not conclusively rule out the possibility that Carter will win, but presumably no one would deny that you can justifiably think or be of the opinion, on the basis of

[16] Hawthorne, Rothschild, and Spectre consider the possibility that the infelicity of (10) and (11) is the result of *neg-raising*, the phenomenon where a negated attribution of belief is interpreted as an attribution of belief in the negation. But they point out that neg-raising is not obligatory since the sentence 'Tim doesn't believe it's raining, nor does he believe it's not raining' sounds fine.

[17] Appealing to similar constructions, Hawthorne, Rothschild, and Spectre suggest that believing even coincides with such attitudes as *suspecting* and *half-expecting*. But we find this data less convincing.

the election polls, that Carter will not win. In the lottery case, you know that you have only strong statistical evidence that ticket 10 did not win, but presumably no one would deny that you can justifiably think or be of the opinion that this ticket did not win.

Turning to the second part of our argument, we will now argue that you would nevertheless be unjustified in believing in the election case that it is not the case that Carter might win and unjustified in believing in the lottery case that it is not the case that ticket 10 might have won. The evidential warrant for believing that $\neg\Diamond\varphi$ is true is considerably higher than the evidential warrant for believing that $\neg\varphi$ is true. When it comes to epistemic modal belief, the slogan "belief is weak" is misleading.

Consider this second cluster of attitude reports:

(14) # Tim is certain that it's not raining, but he doesn't go so far as to believe that it's not the case that it might be raining.[18]

(15) # Tim is free of doubt that it's not raining, but he doesn't go so far as to believe that it's not the case that it might be raining.

(16) # Tim believes that it's not the case that it might be raining, but he isn't certain that it's not raining.

(17) # Tim believes that it's not the case that it might be raining, but he isn't free of doubt that it's not raining.

Like (10)-(13), these sentences are contradictory-sounding.[19] As before, we think that their oddity has a semantic explanation; this linguistic data shows that believing that $\neg\Diamond\varphi$ is true extensionally coincides with such strong epistemic states as being certain and being free of doubt that $\neg\varphi$ is true. If this is right, then we can quickly wrap up our argument for Justification with Risk. In the election case, you are unjustified in being certain or free of doubt, on the basis of the opinion polls, that Carter will not win, so you are unjustified in believing that it is not the case that he might win. In the lottery case, you are unjustified in being certain or free of doubt, on the basis of your purely statistical evidence, that ticket 10 did not win, so you are unjustified in believing that it is not the case that this ticket might have won.

So far we have avoided saying anything precise about what it is to believe that $\neg\Diamond\varphi$ is true. There are, broadly speaking, two different ways

[18] We agree with Stanley [2008] that 'is certain' in sentences like (14) picks out a non-factual *subjective* state.

[19] If you find these sentences hard to parse, feel free to appeal to the duality of *might* and *must* and replace occurrences of 'Tim believes that it's not the case that it might be raining' with 'Tim believes that it must not be raining'. The resulting sentences still sound terrible.

in which philosophers have understood this kind of epistemic modalized belief. The first of these—what Yalcin [2011] calls the "second-order model"—is based on the traditional view of epistemic modals discussed in §3 according to which a speaker who makes a modal claim describes some contextually salient body of information that typically includes her knowledge or potential knowledge state.[20] On this model, believing that $\neg\Diamond\varphi$ is true amounts to believing that this salient information rules out the possibility that φ is true.

A rival "first-order model" of belief involving epistemic modals is endorsed by Yalcin. On this alternative, epistemic modal belief is not a higher-order attitude about what is compatible or incompatible with one's own knowledge or the knowledge of some contextually relevant community of which one is a member. Rather, someone who believes that $\neg\Diamond\varphi$ is true is simply in a first-order doxastic state where the possibility that φ is true is no longer open for them. One believes that $\neg\Diamond\varphi$ is true when one treats the question of whether φ is true as settled in the negative. While evidence suggesting that φ is false can support this belief, one's belief is not *about* this evidence.

Now, we need not decide between these two ways of understanding epistemic modal belief here.[21] We think there is good reason for both first-order and second-order theorists to buy into Justification with Risk. If we understand epistemic modal belief on the second-order model, then in many deliberative contexts where one knows that one's evidence leaves open the possibility that φ is true, one would clearly be unjustified in believing that $\neg\Diamond\varphi$ is true. Perhaps there are some contexts in which one can justifiably believe that a body of information including one's actual

[20] Egan *et al.* [2005] and Dorr and Hawthorne [2013] discuss "shifted uses" of epistemic modals where a speaker describes a body of information that does not involve her actual or potential knowledge. But the existence of such cases makes little difference to our overall argument.

[21] This being said, we think that the first-order model has much to recommend it. First, as Yalcin points out, a question like 'Why believe that it might be raining?' seems equivalent not to the second-order question 'Why believe that I do not know (or we do not know) that it is not raining?' but rather to the question 'Why fail to believe that it is not raining?' This suggests that believing that $\Diamond\varphi$ is true and (by corollary) believing that $\neg\Diamond\varphi$ is true are first-order doxastic states towards the proposition that φ is true. Second, the first-order model is supported by our attitude reports (14)-(17). If being certain and being free of doubt that it's not raining are first-order epistemic states, then second-order theorists must presumably posit linking principles between first-order and second-order states to explain this data. To explain why (14) sounds terrible, for instance, second-order theorists must appeal to a kind of introspection principle requiring that anyone who is certain in a context that $\neg\varphi$ is true must also believe that some salient body of information determined by this context rules out the possibility that φ is true. This principle is controversial.

or easily acquired knowledge rules out some possible state of the world while also knowing that one's evidence leaves open this very possibility. But these contexts will be rare. In any case, presumably no one in our election example has, or can acquire, conclusive evidence prior to the election that Carter will not win, so you cannot justifiably believe that the actual or potential knowledge of *any* agent or group of agents rules out the possibility of a Carter win. Moreover, if we stipulate in our lottery example that you know that no one will have access to the results of the lottery for some time, then you cannot justifiably believe that anyone's actual or potential knowledge before these results are released excludes the possibility that ticket 10 won. In both of our examples, you would be unjustified in forming the relevant negated epistemic possibility belief.

The same holds if we understand modalized belief on the first-order model. In many, if not all, deliberative contexts where one knows that one's evidence leaves open the possibility that φ is true, one would be irresponsible in treating the question of whether φ is true as settled in the negative. To do so would be to willfully ignore the limitations of one's evidence. In some of these contexts, one's evidence might still entitle one to believe that $\neg\varphi$ is true; as we have argued, the evidential warrant for this non-modal belief is lower. One can also provisionally bracket off the possibility that φ is true in hypothetical reasoning contexts. However, when you know that your evidence does not conclusively settle whether φ is true in the negative, you cannot generally be in a first-order doxastic state where the possibility that φ is true is closed.

One might object that believing that $\neg\Diamond\varphi$ is true is justified when the recognizably open possibility that φ is true is extremely remote. To use a well-worn philosophical example, even if your evidence leaves open the possibility that you are a brain in a vat, and you know that it does, aren't you still permitted to believe that it is not the case that you might be a brain in a vat? Well, perhaps you are. But this would not undermine our argument. We have been arguing not for the fully general claim that one would be unjustified in believing that $\neg\Diamond\varphi$ is true in *any* deliberative context where one knows that one's evidence leaves open the possibility that $\neg\varphi$ is true, but only for the weaker claim that the modalized belief would be unjustified in some of these contexts. Note that the possibility of a Carter win and the possibility that ticket 10 won the lottery are not that remote. Though these events are unlikely, it is important not to conflate remoteness with low chance. One does not need to travel very far in logical space to reach worlds in which Carter

wins the election and ticket 10 is chosen.[22]

5 Life without Closure

If our arguments are right, then Single-Premise Closure fails—at least when applied to our ordinary everyday concept of belief exhibited in natural language belief reports. Does this undermine the widespread idea that competent deduction is an epistemically secure means to extend belief? Not entirely.

First, the following non-modal closure principle is immune from our attack:

Single-Premise Closure (Non-Modal Version): For any non-modal sentences φ and ψ, if one is justified in believing that φ is true and one comes to believe that ψ is true on the basis of competently deducing ψ from φ—while justifiably retaining one's belief that φ is true—then one is justified in believing that ψ is true.

So long as φ and ψ do not involve modals, we have not taken issue with the principle that one can responsibly acquire the belief that ψ is true by first responsibly coming to believe that φ is true and then competently inferring ψ from φ.

[22] We have tried to keep our epistemology in this section informal. In a related and complementary paper on epistemic modal belief and closure, Beddor and Goldstein [*ms.*] suggest an attractive way to formalize it. Very briefly, they first model doxastic states using probabilistic structure:

Def. Agent A's *doxastic state* $\langle s_A^w, Pr_A^w \rangle$ in w consists of a privileged domain of worlds $s_A^w \subseteq \mathcal{W}$ and a probability measure Pr_A^w defined over a Boolean algebra of subsets of $\mathcal{W}$ where $Pr_A^w(s_A^w) = 1$. (This resembles the *probability spaces* in Yalcin [2012b])

A Lockean threshold-based semantics for belief (cf. Sturgeon [2008], Foley [2009]) is then formulated in the dynamic semantic framework introduced in n. 9:

$$i[\![Bel_A(\varphi)]\!]_{\mathcal{M}} \quad = \quad i \cap \{w : Pr_A^w(s_A^w[\![\varphi]\!]_{\mathcal{M}}) > t\}$$

That is, A believes that φ is true if the credence assigned to the domain updated with φ is above the threshold t.

Note that on this model, non-modal and epistemic modal belief attributions can report very different features of an agent's doxastic state. If I tell you that Tim believes that it is not raining, I am reporting that Tim assigns sufficiently high credence to the $\neg R$-worlds in his domain. By contrast, if I tell you that Tim believes that it is not the case that it might be raining, I am reporting that none of the worlds in his domain are R-worlds (Beddor and Goldstein are first-order theorists). More needs to be said about what entitles Tim to be in doxastic states with these kinds of structural properties. But even without getting into the details, it should be clear that the evidential warrant for the modalized belief is higher than the evidential warrant for the non-modal belief.

Second, our arguments have been directed at our ordinary concept of belief. Philosophers who talk of 'full' or 'outright' belief often seem to have a more theoretical notion in mind with higher normative standards than ordinary belief (Williamson [2000], Wedgwood [2008]). Depending on how one conceives of this theoretical notion, the following principle might hold:

Single-Premise Closure (Theoretical Version): For any sentences φ and ψ, if one is justified in fully or outright believing that φ is true and one comes to fully or outright believe that ψ is true on the basis of competently deducing ψ from φ—while justifiably retaining one's full or outright belief that φ is true—then one is justified in fully or outright believing that ψ is true.

So closure-lovers who think that a theoretical notion of belief is more important when doing epistemology may not be too concerned by our arguments in this paper. However, when it comes to our basic notion of belief grounded in everyday talk, even the single-premise version of closure must, we think, be abandoned.

References

David Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, 2001.

Bob Beddor and Simon Goldstein. Believing epistemic contradictions. Unpublished manuscript.

Justin Bledin. *Logic Informed*. PhD thesis, University of California, Berkeley, 2013.

Justin Bledin. Logic informed. *Mind*, 123(490):277–316, 2014.

Justin Bledin. *Modus Ponens* defended. Forthcoming in *Journal of Philosophy*.

David Christensen. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford University Press, Oxford, 2004.

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

Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 122(488):867–914, 2013.

Janice Dowell. A flexible contextualist account of epistemic modals. *Philosophers' Imprint*, 11(14):1–25, 2011.

Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Gerhard Preyer and Georg Peter, editors, *Contextualism in Philosophy*, pages 131–169. Oxford University Press, 2005.

Richard Foley. Beliefs, degrees of belief, and the Lockean thesis. In Franz Huber and Christoph Schmidt-Petri, editors, *Degrees of Belief*, pages 37–47. Springer, 2009.

Anthony Gillies. Iffiness. *Semantics & Pragmatics*, 3(4):1–42, 2010.

Ian Hacking. Possibility. *Philosophical Review*, 76(2):143–168, 1967.

John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. Unpublished manuscript.

Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, CVII(3):115–143, 2010.

Angelika Kratzer. The notional category of modality. In Hans-Jürgen Eikmeyer and Hannes Rieser, editors, *Words, Worlds, and Contexts*, pages 38–74. de Gruyter, 1981.

Henry Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, 1961.

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

David C. Makinson. The paradox of the preface. *Analysis*, 25(6):205–207, 1965.

Vann McGee. A counterexample to Modus Ponens. *Journal of Philosophy*, 82(9):462–471, 1985.

Joshua Schechter. Rational self-doubt and the failure of closure. *Philosophical Studies*, 163(2):429–452, 2013.

Robert Stalnaker. Indicative conditionals. *Philosophia*, 5(3):269–286, 1975.

Robert Stalnaker. Assertion. In P. Cole, editor, *Syntax and Semantics 9: Pragmatics*, pages 315–332. Academic Press, 1978.

Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, 1984.

Jason Stanley. Knowledge and certainty. *Philosophical Issues*, 18:33–55, 2008.

William Starr. A uniform theory of conditionals. *Journal of Philosophical Logic*, 43:1019–1064, 2014.

Scott Sturgeon. Reason and the grain of belief. *Noûs*, 42(1):139–165, 2008.

Paul Teller. Epistemic possibility. *Philosophia*, 2(4):303–320, 1972.

Jan Łukasiewicz. Philosophische bermekungen zu mehrwertigen systemen des aussagenkalküls. *Comptes rendus des seances de la Societe des Sciences et des Lettre de Varsovie*, 23, 1930. Translated as 'Philosophical Remarks on Many-Valued Systems of Propositional Logic' in *Selected Works*, pp. 153–178, 1970.

Jan Łukasiewicz. *Selected Works*. North Holland, Amsterdam, 1970.

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

Kai von Fintel and Anthony S. Gillies. 'Might' made right. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 108–130. Oxford University Press, 2011.

Ralph Wedgwood. Contextualism about justified belief. *Philosophers' Imprint*, 8(9):1–20, 2008.

Gregory Wheeler. A review of the lottery paradox. In William Harper and Gregory Wheeler, editors, *Probability and Inference: Essays in Honour of Henry E. Kyburg, Jr.*, pages 108–130. College Publications, 2007.

Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92, 2013.

Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.

Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 295–332. Oxford University Press, 2011.

Seth Yalcin. A counterexample to modus tollens. *Journal of Philosophical Logic*, 41(6):1001–1024, 2012a.

Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, CXII(2):123–160, 2012b.

Igor Yanovich. Standard contextualism strikes back. *Journal of Semantics*, 31(1):67–114, 2014.