Closure and Epistemic Modals*

JUSTIN BLEDIN

Johns Hopkins University

TAMAR LANDO

Columbia University

According to a popular closure principle for epistemic justification, if one is justified in believing each of the premises in set Φ and one comes to believe that ψ on the basis of competently deducing ψ from Φ—while retaining justified beliefs in the premises—then one is justified in believing that ψ. This principle is *prima facie* compelling; it seems to capture the sense in which competent deduction is an epistemically secure means to extend belief. However, even the single-premise version of this closure principle is in conflict with certain seemingly good inferences involving the epistemic possibility modal $\diamondsuit$. According to other compelling principles concerning competent deduction and epistemic justification, one can competently infer $\neg\diamondsuit\varphi$ from $\neg\varphi$ in deliberation even though there are cases in which one can justifiably believe $\neg\varphi$ but would be unjustified in believing $\neg\diamondsuit\varphi$. Thus, as we argue, philosophers must choose between unrestricted closure for justification and the validity of these other principles.

1. Single-Premise 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 of epistemic modals 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 deduce from your justified beliefs. For some φ, though, there are arguably 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 arguably such contexts when *belief* is understood in the ordinary everyday sense. So if philosophers like Yalcin and Veltman are right that inferences from $\neg\varphi$ to $\neg\diamondsuit\varphi$ are valid, then it looks like closure is in trouble. Conversely, if closure for epistemic justification holds, then it

looks like the semantic theories of Veltman and Yalcin must be at least partly rejected.

Keep in mind that it does not just easily fall out of the *meaning* of the expression 'competent deduction' that whenever one competently deduces a conclusion from justified beliefs, one is justified in believing this conclusion. Closure for epistemic justification, after all, is supposed to be a substantive epistemological principle. If it were a simple analytic truth that competently deducing a conclusion from some of your justified beliefs always put you in a position to justifiably believe it, then debates over closure would lose much of their interest.[1]

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:

(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:

(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.[2]

Another class of problematic arguments with the same form concerns past events with low objective chance. Suppose you know 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:

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

---

[1] In fact, *deduction* understood in a more general sense is partially autonomous from belief in that (i) deductions can initiate from assumptions or suppositions rather than beliefs in hypothetical reasoning contexts, and (ii) deductions needn't have any determinate consequences for what one believes (as Harman (1986) stresses, if one deduces a conclusion from some believed premises, then one might come to believe the conclusion, abandon belief in one or more of the premises, or keep things as they are). Of course, the deductions at play in our closure principle do involve expansions of one's belief state.

[2] 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; so they will presumably 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; so (perhaps more controversially) they will presumably also agree that one would be unjustified in believing that it is not the case that Carter might win.

[3] 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 our argument can be made with *either* of our two examples. So those worried about lottery cases are invited to focus on the election case.

But you would be unjustified in believing that:

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

After all, you know that ticket 10 is just as likely to have won the lottery as any other.

Together with preface cases, such lottery cases have sometimes been taken to undermine a *multi-premise* closure under conjunction principle for epistemic justification. It is 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.[4] It is also 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).[5] 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 if, one can competently deduce (2) from (1) or (4) from (3), this, by contrast, conflicts with *single-premise* closure. If there is justification loss in either of our examples, this cannot result from the aggregation of risk across multiple premises.[6]

Note that closure principles for epistemic justification must be sharply distinguished from closure principles for knowledge. According to a naïve closure principle for knowledge, you know what you competently deduce from known premises.[7] One can hold that knowledge is closed under competent deduction without holding that justification is closed under competent deduction, and vice versa.[8] Indeed, one might ultimately come to think that our examples undermine a single-premise closure principle for justification without undermining the analogous principle for knowledge.[9,10]

---

[4] In fact, putting cognitive limitations to the side, you are arguably required 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.

[5] The preface paradox was introduced by Makinson (1965). See Christensen (2004) for a nice discussion of this paradox.

[6] Schechter (2013) also focuses on single-premise closure for epistemic justification. But unlike the issues we are raising here, 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 conclusion given that we are all prone to error in our reasoning.

[7] See Kvanvig (2006) for a critical survey of more refined knowledge principles.

[8] 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.

[9] Bob Beddor and Simon Goldstein (p.c.), on the other hand, argue that one can know on inductive grounds that it will not snow in the Phillipines next year without knowing that it is not the case that it might snow there. If they are right, then the validity of our target inferences would also undermine a single-premise closure principle for knowledge.

[10] See also Lasonen-Aarnio (2008) who argues on the basis of the fallibility of competent deduction (cf. Schechter (2013)) that if a multi-premise closure principle for knowledge fails, then so does the single-premise principle (and for the same kind of reason).

2. A Threefold Tension

More precisely, we have been pointing towards a tension between three principles. The first is single-premise closure for epistemic justification, here formulated with a bit more care:

(5) **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.[11]

The second principle concerns competent deduction (where $\Diamond$ is the epistemic possibility operator):

(6) **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 $\neg\varphi$.[12]

The third principle is this:

(7) **Justification with Risk**

For some φ, there are contexts in which one can justifiably believe $\neg\varphi$ but one would be unjustified in believing $\neg\Diamond\varphi$.

These principles clearly conflict. Consider a sentence φ that witnesses Justification with Risk and suppose that you justifiably believe $\neg\varphi$ in a deliberative context where you would be unjustified in believing $\neg\Diamond\varphi$. Now suppose that you infer $\neg\Diamond\varphi$ from $\neg\varphi$ while retaining your justified belief in the premise. By Łukasiewicz's Principle, you have performed a competent deduction. By Single-Premise Closure, you are justified in believing $\neg\Diamond\varphi$. So something has to give.

Defenders of closure must find a way to reject Łukasiewicz's Principle (and the semantic arguments that stand behind it), Justification with Risk, or both. On the other hand, advocates of Łukasiewicz's Principle who are sympathetic to Justification with Risk must find a way to explain why closure fails for certain inferences involving epistemic modals. Each of these positions, we think, comes with significant costs that are not to be taken lightly. We are not going to take a stand in this paper on which of our three principles should be rejected—not least because the present authors are themselves divided on what the appropriate response to the threefold tension is.

---

[11] Since φ and ψ are schematic variables ranging over sentences, we have been careful to use the natural language truth predicate as a disquotational device at various points in our statement of Single-Premise Closure. But for better readability in the remainder of our paper, we will be loose about use and mention.

[12] Note that our terminology differs slightly 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 deliberative contexts and hypothetical ones triggered by supposition.

Since the motivations for closure are widely acknowledged among philosophers—and because, in addition, we take it as evident that *some* (perhaps restricted) version of that principle is true—we will not make a separate case for Closure. The two remaining principles, Łukasiewicz's Principle and Justification with Risk, are less familiar and therefore more in need of motivation, so we focus below on these other principles. To repeat: *our intention in organizing the paper in this way is not, thereby, to argue for these latter principles and against closure*. Our goal is rather to show that it is far from obvious what the correct choice is.

Before getting started, two clarificatory remarks about Single-Premise Closure are in order. First, which concept of 'justification' enters into this principle? We agree with Schechter (2013) that the relevant notion of justification that supports our closure principle 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)

Second, by quantifying over sentences rather than propositions (as is more commonly done) in our statement of closure, we certainly do not mean to suggest that belief is a relation between thinkers and sentences. We shy away from propositions because, as we will discuss in §3, it is controversial whether sentences with epistemic modals and the indicative conditional express propositions to begin with. Some of the semanticists whose work we will later discuss reject the idea that such sentences are truth-apt constructions expressing propositions. So we want to use a more neutral formulation of Single-Premise Closure—one that allows us to talk about inferences involving epistemic 'might' without prejudging whether 'might'-claims express propositions.

We take this much to be uncontroversial: to believe something is, at least in part, to stand in a relation to a (centered or uncentered) proposition. Moreover, to believe an *unmodalized* φ is to stand in some relation to the proposition that φ. But to have a modalized belief, say, a belief to the effect that it might be raining, may not be to stand in a relation to the proposition that it *might* be raining; for all the arguments of said semanticists show, there may be no such proposition. Rather, to believe that it might be raining may simply be to stand in a relation to the proposition that it *is* raining.[13] So we cannot, without prejudging the issue, say that for any φ, modalized or not, to believe φ is always to stand in a particular relationship to the proposition that φ. For this reason—and because the inferences we are concerned with involve, in an essential way, epistemic modals—we prefer to formulate Single-Premise Closure in terms of sentences rather than propositions.[14,15]

---

[13] For instance, it may be (*inter alia*) for one's belief state to be *compatible* with the proposition that it is raining.

[14] One might object that since closure principles are standardly formulated in terms of propositions rather than sentences, if it turns out that sentences like 'It might be raining' do *not* express propositions, then such sentences cannot figure into counterexamples to Single-Premise Closure. But we think this is too facile; ordinary agents certainly make epistemic modal inferences in their deliberations, and we often report things like 'John believes that it might be raining'. We prefer to formulate Single-Premise Closure in such a way that if it turns out that on the correct semantic theory epistemic modalized sentences do not express propositions, those features of our belief states characterizable using such sentences still fall under this principle.

[15] We are grateful to an anonymous referee for pressing us to clarify this issue.

3. The Case for Łukasiewicz's Principle

We suspect that Łukasiewicz's Principle will meet with a good deal of 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), 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 investigative channels under certain limitations—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, on the standard truth preservation view of validity, 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 modal sentences leaves open the possibility of its truth.

This position, to reiterate, relies heavily on two theses. The first is that sentences involving epistemic modals are truth-apt, and the truth values of these sentences depend on how things stand with respect to contextually supplied bodies of information. The second is that validity should be understood in terms of the preservation of truth at contexts. In recent years, an increasing number of linguists and philosophers have rejected one or both of these theses. Motivated in part by the difficulty in delivering truth conditions for epistemic modal sentences that can make sense of both assertability and disagreement data (see von Fintel and Gillies (2011), MacFarlane (2011), and Yalcin (2011)), these linguists and philosophers have rejected the idea that these sentences even have truth values. Along with this, some have taken the further step of rejecting that idea that valid arguments necessarily preserve truth rather than some other feature of arguments. (If arguments involving epistemic modals can still be evaluated as good or bad, but the sentences they contain are not truth-apt, then the goodness or badness of these arguments cannot have to do with necessary truth-preservation—or so the thinking goes.)

One influential departure from tradition is Veltman (1996). Though he does not explicitly discuss arguments from $\neg\varphi$ to $\neg\diamond\varphi$, Veltman offers a dynamic "update semantics" for epistemic modals and defines various consequence relations over it, each of which validates the Łukasiewicz arguments; his update semantics has since been taken up by Beaver (2001), Willer (2013), Starr (2014), and many others.[16] The basic idea underlying dynamic semantics is that the meaning of a sentence is not its truth conditions but rather a program or instruction for updating information states—what is often called its *context change potential*. Motivated by Stalnaker's (1978) classic view of assertions as proposals to update the common ground of the conversations in which they occur (and earlier work in dynamic semantics by Heim (1982), (1983) and others), Veltman compositionally pairs each sentence in a modal language with a function from information states to information states. On the simplest implementation of this semantics, an information state is just a set of possible worlds—a set of maximally specific

---

[16] For the purposes of this section, we are taking Veltman's dynamic semantics as our starting point, although we could have equally well worked in Yalcin's (2007) static semantic framework.

ways the world might be, or might have been, that are not ruled out by a body of information. So the meaning of a sentence is a function taking sets of worlds to sets of worlds.

To be more precise, it will help to briefly rehearse some of the formal details. Let $\mathcal{L}$ be a sentential modal language containing a countable set $At_{\mathcal{L}}$ of sentence letters A, B,..., the Boolean connectives $\neg$, $\wedge$, and $\vee$, the epistemic possibility modal $\Diamond$, and parentheses. Assume that $\mathcal{L}$ has the usual grammar and let $S_{\mathcal{L}}$ designate the set of sentences in $\mathcal{L}$.

(8) **Update semantics**

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

The *update function* $[\![\,]\!]_{\mathcal{M}} : S_{\mathcal{L}} \to (2^{\mathcal{W}})^{2^{\mathcal{W}}}$ maps each sentence $\varphi \in S_{\mathcal{L}}$ to a function sending each information state $i \in 2^{\mathcal{W}}$ to an information state. (We adopt the conventional notation of representing the result of applying the function $[\![\varphi]\!]_{\mathcal{M}}$ to an information state i by $i[\![\varphi]\!]$ and henceforth leave the relativization to $\mathcal{M}$ implicit.)

$$i[\![A]\!] = \{w \in i : \mathcal{V}(A, w) = 1\}$$
$$i[\![\neg \varphi]\!] = i \setminus i[\![\varphi]\!]^{[17]}$$
$$i[\![\varphi \wedge \psi]\!] = i[\![\varphi]\!] \cap i[\![\psi]\!]^{[18]}$$
$$i[\![\varphi \vee \psi]\!] = i[\![\varphi]\!] \cup i[\![\psi]\!]$$
$$i[\![\Diamond \varphi]\!] = \{w \in i : i[\![\varphi]\!] \neq \emptyset\}^{[19]}$$

The clauses for the sentence letters and Boolean connectives are fairly straightforward. In the atomic case, updating an information state i with A amounts to removing all of the non-A-worlds from i. In the case of negation, updating i with $\neg \varphi$ amounts to first updating i with φ and then removing worlds in the posterior state $i[\![\varphi]\!]$ from the prior state i. In the cases of conjunction and disjunction, updating i with $\varphi \wedge \psi$ and $\varphi \vee \psi$ amounts to first updating i with φ, then updating i with ψ, and finally taking the intersection and union of the posterior states $i[\![\varphi]\!]$ and $i[\![\psi]\!]$ respectively. The last clause is the most interesting. According to Veltman, epistemic possibility imposes a *test* on information states; specifically, $\Diamond \varphi$ tests whether updating i with φ fails to return the empty set $\emptyset$. If this test passes, then the posterior state $i[\![\Diamond \varphi]\!]$ is the input state i itself; otherwise, the output is $\emptyset$. So, for example, $\Diamond A$ tests whether there are any A-worlds in i. If there are, then $i[\![\Diamond A]\!]$ is just i; if there are no A-worlds, then $i[\![\Diamond A]\!]$ is the empty set. (We might also define epistemic necessity $\Box$ and the other Boolean connectives in the usual fashion. Assuming that $\Box$ and $\Diamond$ are duals (interestingly, Veltman (p.c.) himself does not actually think this),

---

[17]   This negation clause appears in Heim (1983).

[18]   In the recent literature on dynamic semantics, the following conjunction clause from Heim (1982) is more common: $i[\![\varphi \wedge \psi]\!] = i[\![\varphi]\!][\![\psi]\!]$. Unlike Veltman's clause (which Beaver (2001) calls "static conjunction"), Heim's clause is non-commutative: if i includes A-worlds and non-A-worlds, then $[\![\Diamond A]\!][\![\neg A]\!] \neq \emptyset$ but $[\![\neg A]\!][\![\Diamond A]\!] = \emptyset$. We were actually a bit surprised to find that Veltman presents a commutative conjunction since some linguists and philosophers of language like Groenendijk and Stokhof (1989) regard a non-commutative conjunction as the essential mark of a dynamic semantics.

[19]   As Yalcin reports in some of his work, the basic idea behind this clause for $\Diamond$ is articulated by Stalnaker (1970).

$\Box\varphi$ tests whether updating i with φ returns i. As before, the output state $i[\![\Box\varphi]\!]$ is i when this test passes and $\emptyset$ when it fails.)

With this semantics in hand, one can then turn to the project of defining semantic consequence. In fact, Veltman defines a few different consequence relations over his dynamic update semantics. We will not go into the details of how these relations differ from one another but it will be useful to take a close look at one of them. Since Veltman's semantics compositionally determines operations on information states (context change potentials) rather than truth conditions, this relation is not defined in terms of unrestricted truth preservation. The basic idea is rather this. Instead of thinking about the contexts in which the premises and conclusion of an argument may or may not be true, we can think about the information states that may or may not *incorporate* (or in Veltman's terminology, "accept") these sentences. And instead of asking "Given a context in which each of the premises are true, is the conclusion also true?", we can ask "Given an information state in which each of the premises is incorporated, is the conclusion also incorporated?" This last question specifies the consequence relation in Veltman (1996) that we focus on.

When does an information state incorporate a sentence? Intuitively: when this information state subsumes the information carried by this sentence.[20] Formally, it is tempting to say this: an information state incorporates a given sentence just in case the sentence is true at every world belonging to the information state. But of course, Veltman cannot say this, because he has not defined a general notion of truth at a world (or more generally, truth at a *point of evaluation* consisting of a world and some other contextual parameters). There is something, however, that he can say, which captures intuitively what we are looking for: an information state incorporates a given sentence just in case this state is immutable under update with this sentence.

(9) **Incorporation**

Information state $i \in 2^{\mathcal{W}}$ *incorporates* $\varphi \in S_{\mathcal{L}}$ (notation: $i \rhd \varphi$) just in case $i[\![\varphi]\!] = i$.

In other words, i incorporates φ just in case i is a fixed point under update with φ. Operate on this information state with the context change potential $[\![\varphi]\!]$ and nothing changes; we get back i again.

Veltman then defines validity in terms of *incorporation preservation*: an argument is valid just in case any information state that incorporates its premises also incorporates its conclusion.

(10) **Informational consequence**[21]

An argument from $\varphi_1,\ldots,\varphi_n$ to ψ is semantically *valid* (notation: $\{\varphi_1,\ldots,\varphi_n\} \vDash \psi$) just in case there is no model $\mathcal{M} = \langle \mathcal{W}, \mathcal{V} \rangle$ such that for some information state $i \in 2^{\mathcal{W}}$, $i \rhd \varphi_1,\ldots,i \rhd \varphi_n$ but $i \not\rhd \psi$. We will also say: ψ is an *informational consequence* of $\{\varphi_1,\ldots,\varphi_n\}$.

---

[20]   *Incorporation* has both formal and informal senses; we must distinguish the technical notion of incorporation that we are about to introduce from the ordinary pre-theoretic sense in which a body of information incorporates, say, that it might be snowing in Canada. Admittedly, we seldom use the expression 'incorporates' in ordinary speech (compare the prevalence of 'is true'). But we do often report incorporation facts using adverbial phrases, e.g., "According to the weather report, it might be snowing in Canada".

[21]   This term is from Yalcin (2007) who defines effectively the same consequence relation over his static semantics.

This way of defining validity captures a sense in which the conclusion of a valid argument is already 'contained' in its premises. If an argument is valid and some information state subsumes the information conveyed by each of its premises, then this state must also subsume the information conveyed by its conclusion. Updating the information state with the conclusion has no effect since the conclusion is already incorporated by this state.

It is worth stressing that informational consequence extensionally coincides with more standard consequence relations defined in terms of truth preservation over simple fragments of language without epistemic modals and the indicative conditional—such as the fragment in which almost all mathematics is formulated. It is only when we start looking at the kind of modal and conditional inferences discussed in this paper that these consequence relations come apart and we must choose between them.

It is easy to verify that our target Łukasiewicz arguments come out valid on the above definitions: $\{\neg\varphi\} \vDash \neg\Diamond\varphi$.[22] If $i \vartriangleright \neg\varphi$, then $i[\![\neg\varphi]\!] = i$, so $i[\![\varphi]\!] = \emptyset$. But then $i[\![\Diamond\varphi]\!] = \emptyset$, so $i[\![\neg\Diamond\varphi]\!] = i$ and $i \vartriangleright \neg \Diamond\varphi$. Information incorporating $\neg\varphi$ must also incorporate $\neg\Diamond\varphi$, so arguments like those in our election and lottery examples are (on these definitions) semantically valid. Of course, Łukasiewicz's Principle states not that arguments from $\neg\varphi$ to $\neg\Diamond\varphi$ are valid but rather that one can competently make such inferences in both categorical and hypothetical contexts. But we think that it is a fairly easy step from validity to competent inference. Indeed, it would seem odd to on the one hand accept that an argument form is valid and on the other hand think that this inference is unreliable in deliberation. While some philosophers might think that one can competently make certain inferences that are not semantically valid (cf. Stalnaker (1975) on *reasonable inference*; more on this below), we are not aware of anyone who thinks that one cannot competently make valid inferences in deliberation.[23]

We have shown at this point that one can provide a semantics for a modal language that (i) has at least some intuitive appeal, (ii) does not assign truth conditions to epistemic modal sentences, and (iii) validates Łukasiewicz arguments. What more can be said in its favor? Let us now explore some reasons why one might prefer the kind of informational account in Veltman (1996) to its more traditional truth-centric rivals.

One very general motivating reason is, as briefly mentioned above, pessimism about the prospects of assigning satisfactory truth conditions to epistemic modal sentences. If one thinks that such sentences are not truth-apt but they can nevertheless enter into valid argument patterns, then one will need a consequence relation which, like informational consequence, does not require truth preservation. Standard conceptions of validity in terms of truth preservation will not do for arguments involving sentences that do not have truth values.[24]

Alternatively, one might be entirely unmoved by the thorny problems involved in assigning truth values to modalized sentences but simply be disappointed at the verdicts

---

[22] To reiterate, Łukasiewicz arguments are validated not just by informational consequence; each of the consequence relations defined by Veltman over his update semantics validates arguments of this form.

[23] Note that informational consequence also validates inferences from φ to $\Box\varphi$. So, as an anonymous referee asks, why not work with this simpler argument form and avoid the need to discuss negation? Well, admittedly, inferences from φ to $\Box\varphi$ are slightly odd. While we think that the perceived oddness stems from their redundancy (compare reiteration of φ) not from their unreliability, we prefer to base our paper around the more compelling Łukasiewicz arguments.

[24] Another general motivating reason for going informational that we do not have space to discuss in more detail is that it allows one to develop a unified semantics and logic for declarative and interrogative sentences. See work in *inquisitive semantics* by Ciardelli, Groenendijk, and Roelofsen (2015) and Ciardelli (2016), [*ms.*].

about validity that many truth conditional accounts issue in specific cases. One might therefore be driven to look for an alternative semantic account that better honors our discriminating judgments with respect to various argument forms—a semantic account that validates intuitively good arguments that come out invalid on other accounts, and invalidates intuitively bad arguments that come out valid on other accounts.

Some philosophers have argued that informational consequence does well in this regard.[25] Consider the argument from the premises

(11) Either the price of honey will drop or the demand for sorghum will rise.

(12) The price of honey might not drop.

to the conclusion

(13) The demand for sorghum might rise.

This argument seems intuitively good. But on standard contextualist semantics for epistemic modals, it fails to preserve truth. (Imagine that the price of honey will drop in the context of use, but the speaker's knowledge, or some group's knowledge, or some other salient information state supplied by this context leaves open the possibility that it will not. Add to this that this relevant information state rules out the possibility that the demand for sorghum will increase.) By contrast, this argument is validated by informational consequence. If some (nonempty) information state i incorporates (11), then *each* world in i is either one in which the price of honey will drop or one in which the demand for sorghum will rise (or both). If i also incorporates (12) then this state must include a world in which the price of honey will not drop. But this world must then be one in which the demand for sorghum rises, so i incorporates (13) as desired.

For another example, consider the argument from the premises

(14) Either the supply of corn or soybeans will rise.

(15) The supply of corn will not rise.

to the conclusion

(16) It must be the case that the supply of soybeans will rise.

(Important: the 'must' here should be read epistemically.) Again, this argument seems intuitively good. But it comes out invalid on standard contextualist semantics. (Imagine that the supply of soybeans will rise and the supply of corn will fall or stagnate in the context of use, but the information state supplied by this context leaves open the possibility that the supply of soybeans will not rise.) By contrast, this argument is validated by informational consequence. If i incorporates (14), then each world in i is either one in which the supply of corn will rise or one in which the supply of soybeans will rise (or both). If i also incorporates (15), then in each of these worlds, the supply of corn will not rise. But then each world in i is one in which the supply of soybeans will rise, so i incorporates (16) as desired.

---

[25] In fact, Yalcin (2007) and Bledin (2014) appeal to the goodness of Łukasiewicz arguments themselves in arguing for informational consequence. But to do so here would be to beg the question.

Adding the indicative conditional to our language $\mathcal{L}$ provides further examples of argument forms that defenders of informational consequence have thought to play in their favor.[26] The most famous of these forms is *modus ponens* (we have already provided a sample MP argument in n. 2 above).[27] As Khoo (2013) discusses, neither Kratzer's (1986), (2012) influential truth conditional semantics for conditionals nor Gillies' (2009) semantics validates *modus ponens*.[28] Neither does the formal semantics in Kolodny and MacFarlane (2010) to take another example.[29] But *modus ponens* is valid on the informational view (at least when combined with a standard dynamic entry for the indicative).

To see this, let us supplement Veltman's update semantics in (8) with the following semantic clause for the indicative conditional $\Rightarrow$ due to Gillies (2010):

(17) **Update semantics for the indicative**

$i[\![\varphi \Rightarrow \psi]\!] = \{w \in i : i[\![\varphi]\!][\![\psi]\!] = i[\![\varphi]\!]\}$[30]

Like epistemic modals, an indicative conditional $\varphi \Rightarrow \psi$ imposes a test on information states. If sequentially updating i with its antecedent φ followed by its consequent ψ delivers the same output as updating i with φ alone, then this test returns i; otherwise, it returns $\emptyset$. It is easy to see that *modus ponens* must then come out valid: if $i \vartriangleright \varphi$, then the test condition for $\varphi \Rightarrow \psi$ becomes $i[\![\psi]\!] = i$, so $i \vartriangleright \varphi \Rightarrow \psi$ only if $i \vartriangleright \psi$.

Now, none of the data points considered so far are decisive.[31] Most importantly, opponents of informational consequence might agree that the above arguments are good ones but nevertheless insist that they are not *semantically valid* by allowing the concepts of validity and good argument to extensionally come apart. Khoo (2013), for one, suggests that even though Kratzer and Gillies invalidate *modus ponens* for the indicative, they might still call *modus ponens* arguments "dynamically valid" since in any context in which we are asserting or

---

[26]   Adding quantifiers generates more examples. See Yalcin (2015) for discussion.

[27]   While *modus ponens* for the indicative conditional has perennially come under attack (by McGee (1985), Kolodny and MacFarlane (2010), Willer (2010), and others), Bledin (2014), (2015) defends the reliability of this rule.

[28]   At least not when combined with a standard consequence relation that preserves truth at a context. Khoo's discussion is a bit misleading since Gillies (2010) argues for an alternative "update-to-test" entailment relation where the conclusion of an argument is evaluated against a context incorporating the information carried by its premises (this is essentially one of the other consequence relations in Veltman (1996) that we did not discuss).

[29]   Kolodny and MacFarlane call *modus ponens* arguments "quasi-valid" since these arguments are reliable in categorical contexts where the premises are known. But they argue that *modus ponens* is unreliable in hypothetical contexts.

[30]   Heim (1983) presents this alternative clause: $i[\![\varphi \Rightarrow \psi]\!] = i \setminus (i[\![\varphi]\!] \setminus i[\![\varphi]\!][\![\psi]\!])$. Unlike Gillies' test semantic clause, updating i with $\varphi \Rightarrow \psi$ here amounts to first updating i with φ, then updating the resulting state $i[\![\varphi]\!]$ with ψ, and finally removing the worlds in $i[\![\varphi]\!]$ that are not also in $i[\![\varphi]\!][\![\psi]\!]$ from the original state i. We could have equally well worked with this Heimian entry since, combined with informational consequence, it also upholds *modus ponens*.

[31]   While we have considered only intuitively good arguments that informational consequence validates, another source of support for this relation are intuitively bad arguments and pieces of argumentation. Yalcin (2012) presents several examples of intuitively bad *modus tollens* arguments that informational consequence invalidates. (Some might worry about the asymmetry—shouldn't a viable consequence relation validate (or invalidate) both *modus ponens* and *modus tollens*? But Yalcin argues that the asymmetry is in fact well-grounded.) Bledin (2014), (2015) also develops an informational account of deductive inquiry based on informational consequence that aims to predict and explain the badness we sometimes find when applying *reductio* and proof by cases in languages with epistemic modals and the indicative conditional (see Kolodny and MacFarlane (2010) for nice examples of this badness).

supposing that an indicative conditional and its antecedent are true, its consequent holds; here he echoes Stalnaker's (1975) well-known distinction between the semantic concept of *entailment* and the pragmatic concept of *reasonable inference*.[32] Extending the suggestion to the other arguments in this section, one might think also that the premises of those arguments do not actually *entail* their conclusions but these arguments are still *reasonable*.

To flesh this out a bit more, take the standard contextualist, or for that matter anyone who holds that validity is a matter of necessary truth-preservation. Such a philosopher might ask herself two very *different* questions. First, which arguments are such that if the premises are true, then the conclusion must also be true? And, second, which arguments are such that an agent who accepts or takes on board the premises inside of an episode of deliberation is reasonable in accepting the conclusion? It may well be that with respect to a particular argument, the answer to the first question is 'No' (the argument is not truth-preserving) whereas the answer to the second question is 'Yes'.

When it comes to our inferences from $\neg\varphi$ to $\neg\Diamond\varphi$, in particular, it is easy to see how the answers to these questions can come apart. If you are a contextualist, for example, then an argument of this form might fail to preserve truth because, as we have already discussed, in some context where φ is false, the salient body of information does not rule out φ. So the premise is true but the conclusion is false. Nevertheless, in taking on board the premise $\neg\varphi$ inside of an episode of deliberation, the reasoner allows this premise to 'shape' her own information state—the information state in relation to which she performs her reasoning. And once she takes on board $\neg\varphi$, she can no longer treat the question of whether φ as open. She has already shaped her information state to exclude φ. So a philosopher of this stripe can concede that while the argument is not truth-preserving, an agent that takes on board the premise can, inside of that episode of deliberation, competently deduce the conclusion. Again, this has everything to do with the dynamics of deliberation and very little to do with truth-preservation.

Now, on the more general matter of whether we should admit this kind of separation between validity and deductively good argument, the authors of this paper are somewhat divided. One of the authors is inclined to think that the distinction between entailment and reasonable inference, to its credit, allows us to understand validity in the usual way, in terms of truth-preservation, while explaining what seems 'good' about the Łukasiewicz inferences, as well as some of the other inferences highlighted in this section.

Meanwhile, the other author finds the proposed disconnect between validity and deductive goodness unattractive, especially when we have a serious rival to truth preservation, informational consequence, that fits nicely with the informational picture of deliberation sketched in the previous paragraph, and accommodates arguments involving non-truth-apt sentences. Keeping in mind that the informational notion of validity coincides with necessary truth preservation over the fragment of $\mathcal{L}$ without modals and the indicative conditional, this authors wonders why we should cling to truth preservation.

But for present purposes, this disagreement does not matter much. Note that appealing to the distinction between reasonable inference and validity still does not dissolve the

---

[32]   Stalnaker argues, for example, that 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 it can fail to preserve truth, but this inference is nevertheless reasonable since if one can appropriately assert the disjunctive premise in any context of use, then the background common ground information in this context will incorporate the conditional conclusion.

tension with Single-Premise Closure. The philosopher who wishes to argue against closure does not actually need the strong claim that Łukasiewicz arguments are *valid*; she needs only the weaker claim that one can competently deduce the conclusion of a Łukasiewicz argument from its premise in the course of deliberation. So if Łukasiewicz arguments are reasonable in this way, there is still a conflict with Single-Premise Closure and Justification with Risk.

In conclusion, the dynamic update semantics defended by Veltman and others is controversial but it provides an interesting alternative to standard contextualist semantics of epistemic modality—one that predicts the intuitive goodness of Łukasiewicz arguments. Whether one thinks the goodness of these inferences has to do with formal validity, or whether one reserves 'validity' for necessary truth-preservation, the discussion in this section shows, we think, that it is far from obvious whether Łukasiewicz's Principle should be abandoned.

4. The Case for Justification with Risk

What about Justification with Risk? Moving on to this principle, we will now present a case that there are contexts in which you are in fact justified in believing that $\neg\varphi$ is true but would be unjustified in believing that $\neg\Diamond\varphi$ is true—at least such contexts exist when working with our ordinary concept of belief. Indeed, our two examples from §1 arguably fit this mold. In the election case, we will now argue, 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 likewise 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 case for Justification with Risk has two parts. First, we argue that there are contexts in which one can justifiably believe $\neg\varphi$ even though one's evidence leaves open the possibility that φ is true and one knows that it does. Second, we argue that in at least some of these contexts, one would be unjustified in believing [φ]. Putting these two pieces together, the argument will be complete.

Since we are dealing with our ordinary 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 (2016):

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

(19) # 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, it is not clear how this would go.[33] 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:

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

(21) # 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 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.[34]

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 such pre-theoretically weak states as thinking or being of the opinion that φ is true.[35] If this is right, then 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 form beliefs, there are presumably contexts in which one can justifiably believe that $\neg\varphi$ is true despite knowing that one's evidence leaves open the possibility that φ is true (at least when φ is non-modal). After all, if one has recognizably non-conclusive evidence that overwhelmingly supports that $\neg\varphi$ is true, then presumably one can still justifiably *think* or *be of the opinion* that $\neg\varphi$ is true. In the election case, you know that your available evidence does not conclusively rule out the possibility that Carter will win, but you can still justifiably think or be of the opinion, on the basis of 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 the lottery, but you can still justifiably think or be of the opinion that this ticket did not win.

To summarize, the first part of our case for Justification with Risk combines the following two theses:

(22) For any non-modal sentence φ, one believes that φ is true if and only if one thinks/is of the opinion that φ is true.

(23) For some non-modal sentence φ, there are contexts in which one can justifiably think/be of the opinion that $\neg\varphi$ is true even though one's evidence leaves open the possibility that φ is true and one knows that it does.

The first thesis explains the infelicity of (18)–(21). The second thesis is supported by ordinary intuitions about what it takes to responsibly think or be of the opinion that something is the case. Admittedly, the argument here is not decisive. Some philosophers might reject (23) and deny, say, that being of the opinion that φ is true is weak in the sense that Hawthorne *et al.* have in mind. For instance, proponents of a knowledge norm for belief (e.g., Williamson (2000) and Sutton (2005)) might object to our analysis of the lottery case as follows: "Look, if someone has purely statistical evidence about the lottery, then they're not justified in being of the opinion that ticket 10 did not win. They're really only justified in being of the opinion that it is extremely unlikely that this ticket won. One can be of the opinion that φ is true only if one *knows* that φ is true. To be of the opinion that ticket 10 did not win the lottery is to be over-opinionated."[36]

---

[34] Appealing to similar constructions, Hawthorne, Rothschild, and Spectre suggest that believing even coincides with such attitudes as *suspecting* and *half-expecting*. But we do not find these data as convincing.

[35] Hawthorne, Rothschild, and Spectre concede that there might be a theoretical notion of belief with a higher evidential standard. A bit more on this in §5.

[36] We are grateful to John MacFarlane (p.c.) for pressing this response.

However, this line of resistance is hard to square with ordinary speech. Consider the following exchange:

(24) A: If ticket 10 won, I can pay off my credit card debt.
    B: What are the chances?
    A: One in a thousand.
    B: Well then I really think you're not going to be paying off your credit card debt anytime soon. In my opinion, your ticket lost.
    A: ?But you haven't heard the results of the lottery. For all you know, my ticket won.[37]
    B: That's right, I haven't. But I didn't say that I'm *sure* ticket 10 lost. I'm only of the opinion that it lost on the basis of the statistics and that's good enough for me.

To our ears at least, A's last reply is odd and B's counter-reply is fine. This is the opposite of what we should expect if being of the opinion is governed by a high standard like knowledge. If the justificatory standard for being of the opinion is knowledge, then A's reply is a legitimate attack on B's opinion report and B's counter-reply is an admission of epistemic irrationality.[38]

Turning to the second part of the case for Justification for Risk, we will now argue that you would 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. More generally, we will argue that the evidential warrant for believing $\neg\Diamond\varphi$ is considerably higher than the evidential warrant for believing $\neg\varphi$. At least when it comes to epistemic modal belief, the slogan "belief is weak" is misleading.[39]

Why?[40] So far we have avoided saying anything precise about what it is to have this kind of epistemic modalized belief. There are, broadly speaking, two different ways in which philosophers have understood it. The first of these—what Yalcin (2011) calls the "second-order model"—is based on the traditional contextualist view of epistemic modals

---

[37] Sentences marked with '?' sound a bit better to us than those marked with '#'.

[38] Some philosophers might try to explain the oddness of A's reply in (24) by arguing that B's opinion report should not be taken literally; B does not actually express that she is of the opinion that φ is true but only that she is of the opinion that φ is *probably* true (Sutton (2005) and Stanley (2008) suggest this reading of some ordinary language belief reports). However, we agree with Hawthorne, Rothschild, and Spectre (2016) that positing this kind of mismatch should be avoided if possible. Moreover, B's follow-up reply is strange if opinion reports are non-literal.

[39] To be clear: when we say that the evidential warrant for believing $\neg\Diamond\varphi$ is high, we have in mind the strength of evidence needed regarding the falsity of φ, and not regarding the truth or falsity of some modalized claim.

[40] Some preliminary cursory support for this thesis comes from this second cluster of attitude reports:

(i) # Tim is certain that it's not raining, but he doesn't believe that it's not the case that it might be raining. (We agree with Stanley (2008) that the 'is certain' here picks out a non-factive subjective state.)
(ii) # Tim is free of doubt that it's not raining, but he doesn't believe that it's not the case that it might be raining.
(iii) # Tim believes that it's not the case that it might be raining, but he isn't certain that it's not raining.
(iv) # 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 (18)–(21), these sentences are contradictory-sounding. As before, we think that their oddity has a semantic explanation; these linguistic data show 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 the case for Justification with Risk can be quickly wrapped up. 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 presumably 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 presumably you are unjustified in believing that it is not the case that this ticket might have won.

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.[41] On this model, believing $\neg\Diamond\varphi$ 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 modalized 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, to believe $\neg\Diamond\varphi$ is to take up an attitude towards φ itself. In particular, someone who believes $\neg\Diamond\varphi$ is simply in a first-order doxastic state where the possibility that φ is true is no longer treated as open. One believes $\neg\Diamond\varphi$ 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.[42]

Now, we will not decide here between these two different ways of understanding epistemic modalized belief. We bring them up only to point out that there is good reason for both first-order and second-order theorists alike to buy into the second part of the case for Justification with Risk.

If we understand epistemic modal belief on the second-order model, then in many, if not all, contexts in which one knows that one's evidence leaves open the possibility that φ is true, one would surely 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 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 even acquire, conclusive evidence prior to the election being held that Carter will not win, so you cannot justifiably believe that the actual or potential knowledge state of *any* individual or group of individuals 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 at a time 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, 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. Someone 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 the second part of the case for Justification with Risk. We have been arguing not for the fully general claim that

---

[41]   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 fringe cases makes little difference to our overall argument here.

[42]   Yalcin actually presents a more sophisticated account where one believes $\sim\Diamond\varphi$ when one's belief state is sensitive to the question of whether $[\varphi]$ and this state rules out the possibility that $[\phi]$ is true with respect to this question. But let us ignore this extra complexity here.

for *any* φ, one is unjustified in believing $\neg\Diamond\varphi$ if one knows that one's evidence leaves open the possibility that $\neg\varphi$ is true, but only for the much weaker claim that for *some* φ, in *some* deliberative context, the modalized belief is unjustified. 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 or ones in which ticket 10 is chosen.[43]

5. Conclusion

Where does this leave us? Someone who wants to hold onto Single-Premise Closure must reject either Łukasiewicz's Principle or Justification with Risk. If they go the first route, they must not only reject the claim that Łukasiewicz's arguments are valid, but also the weaker claim that we can competently make these arguments inside an episode of deliberation. If they go the second route, then in each of our sample arguments, they must either deny that one can justifiably believe the premise, e.g., that ticket 10 in the lottery did not win, or hold that one can in fact justifiably believe the conclusion, e.g., that it is not the case that ticket 10 might have won. If one thinks that the premise belief is unjustified, this might be because one thinks that the linguistic data presented in §4 does not establish that belief is weak; alternatively, it might be because one has a more theoretical notion of 'full' or 'outright' belief in mind with higher justificatory standards than the ordinary, everyday concept (Williamson (2000), Wedgwood (2008)) and it is this theoretical notion that enters into closure for epistemic justification.

Someone convinced by the cases presented in §3 and §4, on the other hand, must give up Single-Premise Closure. They might still be able to hold onto a restricted version of this principle that applies only to a language devoid of epistemic modals and the indicative conditional:

(25) **Non-Modal Single-Premise Closure**

> For *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.

---

[43] We have kept our epistemology in this section informal. In a related paper on epistemic modal belief and closure, Beddor and Goldstein [*ms.*] present a formal epistemology on which Justification with Risk holds. Very briefly, they first model doxastic states using probabilistic structure:

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 $Pr_A^w(s_A^w) = 1$.

A threshold semantics for belief (cf. Sturgeon (2008), Foley (2009)) is then formulated in the update semantic framework in §3: $i[\![Bel_A(\varphi)]\!] = i \cap \{w : Pr_A^w(s_A^w[\![\varphi]\!]) > t\}$. 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 non-raining-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 non-raining-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 features. But even without getting into the details, it should be clear that the evidential warrant for the modal belief is higher than the evidential warrant for the non-modal belief.

This restricted principle is immune from the challenges raised above. One of the more interesting upshots of our discussion, we think, is that defenders of the kind of informational account sketched in §3 who also accept Justification with Risk call for an even more radical departure from orthodoxy than previously recognized. Not only do they revise core semantic notions like validity and consequence, but they also call into question commonly held epistemological principles like closure for justification.

References

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

Bob Beddor and Simon Goldstein. Believing epistemic contradictions. *Review of Symbolic Logic*, Forthcoming.

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

——. *Modus Ponens* defended. *Journal of Philosophy*, CXII(2):57–83, 2015.

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

Ivano Ciardelli. *Questions in Logic*. PhD thesis, University of Amsterdam, 2016.

——. Questions as information types. Unpublished manuscript.

——, Jeroen Groenendijk, and Floris Roelofsen. *Inquisitive Semantics*. ESSLLI 2015 Lecture notes.

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. On truth-conditions for If (but not quite only If). *Philosophical Review*, 118(3):325–349, 2009.

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

Jeroen Groenendijk and Martin Stokhof. Context and information in dynamic semantics. In Klaus von Huesinger and Urs Egli, editors, *Working Models of Human Perception*, pages 457–486. Academic Press, 1989.

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

Gilbert Harman. *Change in View: Principles of Reasoning*. MIT Press, Cambridge, 1986.

John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173(5):1393–1404, 2016.

Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, 1982.

——. On the projection problem for presuppositions. In Michael Barlow, Daniel Flickinger, and Michael Wescoat, editors, *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*, pages 114–125. Stanford University, 1983.

Justin Khoo. A note on Gibbard's proof. *Philosophical Studies*, 166(1):153–164, 2013.

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.

——. Conditionals. *Chicago Linguistics Society*, 22(2):1–15, 1986.

——. *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press, Oxford, 2012.

Jonathan Kvanvig. Closure principles. *Philosophy Compass*, 1(3):256–267, 2006.

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

Maria Lasonen-Aarnio. Single premise deduction and risk. *Philosophical Studies*, 141(2):157–173, 2008.

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

John MacFarlane. Epistemic modals are assessment-sensitive. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 144–178. Oxford University Press, 2011.

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

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

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

Robert Stalnaker. Pragmatics. *Synthese*, 22(1–2):272–289, 1970.

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

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

——. *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.

Jonathan Sutton. Stick to what you know. *Noûs*, 39(3):359–396, 2005.

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.

——. *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. New surprises for the Ramsey test. *Synthese*, 176(2):291–309, 2010.

——. 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.

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

——. A counterexample to Modus Tollens. *Journal of Philosophical Logic*, 41(6):1001–1024, 2012.

——. Epistemic modality *De Re*. *Ergo*, 2(19):475–527, 2015.

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