Modal Knowledge for Expressivists

Preprint – forthcoming in *Journal of Philosophical Logic*

Peter Hawke*

Abstract

What does 'Smith knows that it might be raining' mean? Expressivism here faces a challenge, as its basic forms entail a pernicious type of *transparency*, according to which 'Smith knows that it might be raining' is equivalent to 'it is consistent with everything that Smith knows that it is raining' or 'Smith doesn't know that it isn't raining'. Pernicious transparency has direct counterexamples and undermines vanilla principles of epistemic logic, such as that knowledge entails true belief and that something can be true without one knowing it might be. I re-frame the challenge in precise terms and propose a novel expressivist formal semantics that meets it by exploiting (i) the topic-sensitivity and fragmentation of knowledge and belief states and (ii) the apparent context-sensitivity of epistemic modality. The resulting form of *assertibility semantics* advances the state of the art for state-based bilateral semantics by combining attitude reports with context-sensitive modal claims, while evading various objectionable features. In appendices, I compare the proposed system to Beddor and Goldstein's 'safety semantics' and discuss its analysis of a modal Gettier case due to Moss.

Keywords: expressivism; epistemic possibility; knowledge attribution; contextualism; formal semantics; fragmentation; subject matter

1 Introduction

Natural language has nuanced resources for signaling an agent's epistemic position. Compare:

(1) It isn't raining.

(2) It might be raining.

We take (1) to be a straightforward *description*: a declarative sentence with the canonical discourse role of (i) representing the world as being a certain way, (ii) signaling that the speaker is committed to the world being that way, and (iii) inviting

*Department of Philosophy and Hong Kong Catastrophic Risk Centre, Lingnan University, Hong Kong, peterhawke@ln.edu.hk.

interlocutors to share this commitment. We call sentences like (2) *bare might claims*: a declarative formed by applying the 'might' operator to straightforward description p. (We write *might-p* as shorthand for 'it might be that p'.) The modal in (2) invites an epistemic reading. Intuitively, (2) is aptly asserted when the available information doesn't establish that it isn't raining: someone who *knows* it isn't raining aptly asserts (1), not (2). Throughout, we take 'might' (and duals 'must' and 'can't') to have an ordinary epistemic reading, whatever exactly it is.

What is the canonical discourse role of bare might claims? According to *expressivism*, (2) isn't a description: the characteristic job of might-p is *not* merely to propagate commitment to a representational content.[1] Indeed, one paradigmatic version of expressivism says that might-p has *no* representational content; so uttering it is not to, for example, describe oneself as lacking information that rules out p. Just as asserting description p *expresses* that one accepts things are as p describes without *asserting* one accepts this, so asserting might-p expresses that one does *not* accept things are *not* as p describes, without asserting one doesn't accept this.[2]

Influential motivation for expressivism comes from the impression that might-p defies the *logic* of description.[3] Intuitively, (1) and (2) are opposed: if Ann says (1) and Bob replies (2), they disagree; asserting both (1) and (2) sounds incoherent, with the incoherence surviving embedding ('Ann thinks it isn't raining and might be raining' and 'if it isn't raining and might be raining, we need an umbrella' sound jarring). Now, if not-p and might-p were *contradictory descriptions*, it would follow (assuming their semantic presuppositions are met) that might-p entails p. But 'it might be raining' does *not* entail 'it is raining': one can commit to the former without the latter, without discernible presupposition failure. Expressivists instead posit that (1) and (2) together issue an 'expressivistic' contradiction (they simultaneously express acceptance and *lack* of acceptance that it isn't raining), despite commitment to (2) not implying commitment to 'it is raining' (aptly asserting the former only requires not accepting that it isn't raining; the latter requires accepting that it *is* raining).

But an unresolved issue is whether expressivists can deliver a plausible semantics for *attitude ascriptions* that embed a bare might claim. As many observe,[4] *crude* forms of expressivism (as we label them) cannot meet this challenge, in virtue of exhibiting a property we call *pernicious transparency*. Pernicious transparency implies that 'Smith knows it might be raining' is equivalent to 'Smith doesn't know that it isn't raining' or, alternatively, 'it is consistent with the sum total of Smith's knowledge that it is raining' (typically glossed as 'for all Smith knows, it is raining'). Likewise, pernicious transparency implies that 'Smith believes it might be raining' is equivalent to 'Smith doesn't believe that it isn't raining' or, alternatively, 'it is consistent with the sum

[1] Characterizing expressivism pragmatically follows, for instance, Yalcin [2011] and [Incurvati and Schlöder, 2023, Ch.1].

[2] Compare Yalcin [2007, 2011, 2012a], MacFarlane [2014, Sect. 10], Willer [2013, 2017a], Cantwell [2014, 2021], Hawke and Steinert-Threlkeld [2015, 2018, 2021], Lin [2016], Incurvati and Schlöder [2019, 2023]. Related accounts include Veltman [1996], Moss [2015].

[3] See Yalcin [2007] and Bledin [2014].

[4] See [Yalcin, 2011, Sect.5], [Beddor and Goldstein, 2021, Sect. II].

total of Smith's beliefs that it is raining'. But this predicts a logic of knowledge ascription that defies intuition and theoretical orthodoxy. Most basically, it issues counter-intuitive predictions about equivalency. Compare:

(3) # Bob isn't here but Ann knows Bob might be here.

(4) Bob isn't here but Ann doesn't know he isn't here.

(5) Bob isn't here but, for all Ann knows, he is.

Pernicious transparency says (3) is equivalent to (4) or (5). But far from conversational uniformity, an out-of-the-blue assertion of (3) can be markedly odd and hard to interpret (compare the benign 'Bob isn't here but Ann mistakenly believes he might be'), while (4) and (5) are straightforward declarations of Ann's ignorance.

Further, as the literature notes, pernicious transparency conflicts with orthodox principles of epistemic logic: that knowledge is factive, that knowledge entails belief, and that something can be true without one knowing it might be.[5] (§2.2 spells this out.)

Hence, *crude* expressivism is widely rejected. So, expressivists face a challenge. The semantics of Yalcin [2007] – a well-spring of recent interest in expressivism for might-p – exemplifies crude expressivism. Critics of expressivism thus bolster their case with the dire predictions of pernicious transparency [Dorr and Hawthorne, 2013, Holliday and Mandelkern, forthcoming], while friends of expressivism grapple with what crude expressivism should be replaced with [Yalcin, 2012b, Beddor and Goldstein, 2021].

> **The Challenge From Pernicious Transparency:** can an expressivist theory avoid pernicious transparency and the counterintuitive/unorthodox epistemic logic it implies (or at least convincingly motivate specific departures from intuition/orthodoxy), while preserving whatever clear-cut advantages crude expressivism has over descriptivist rivals?

The primary goal of this paper is to *sharpen* and *answer* this challenge: we reformulate the challenge in a precise formal setting, then propose a novel expressivist semantics that demonstrably meets it. Pernicious transparency and its dire consequences for epistemic logic are hereby evaded, without abandoning key expressivist advantages. Notably, the sharpened challenge extends to *descriptivists*: once it is clarified that crude expressivism yields attractive principles for epistemic logic that straightforward descriptivism *lacks*, descriptivists are likewise challenged to explain the *prima facie* appeal of these principles without embracing pernicious transparency.

Our answer to the challenge is two-pronged. Pernicious transparency, we argue, is closely related to *Holism*, a familiar but unpopular account of the structure and attribution of knowledge and belief.[6] Despite admiration as an idealized modeling tool, Holism finds little endorsement as a realistic picture of epistemic states. Meanwhile,

[5] See Yalcin [2012b], Dorr and Hawthorne [2013, fn.7], Moss [2018, §6.1], Beddor and Goldstein [2021].

[6] Exemplified by the influential treatment of epistemic logic by Hintikka [1962].

if judiciously implemented, standard tools for refinement – *topic/question-sensitivity*[7] and *fragmentation*[8] – allow expressivists to evade pernicious transparency and preserve principles like: knowledge entails belief. That's the first prong.

However, fragmentation and topic-sensitivity, as we implement them, do not by themselves secure factivity for claims of modal knowledge. Further, they permit a *restricted* but still problematic form of transparency. The paper develops a nuanced response. First, I argue that the preliminary linguistic data is *equivocal* on whether modal knowledge is factive. Second, I show that expressivists can account for the puzzling data: supplementing our topic-sensitive and fragmented expressivism with an independently motivated form of context-sensitivity predicts that attribution of modal knowledge exhibits factivity in *some* prominent contexts, but not all. As a bonus, this system evades even restricted transparency. This is the second prong.

We develop our theory in the paradigm of formal *assertibility semantics* (or 'acceptance semantics'), a promising perspective for capturing the core structure of expressivism for epistemic modals.[9] More generally, *bilateral state-based systems* are growing in stature in formal semantics.[10] We advance the state of the art with novel systems that combine attitude reports with context-sensitive modal claims, built on a foundation familiar from Yang and Väänänen [2017], Hawke and Steinert-Threlkeld [2021], Ciardelli [2021], and Aloni [2022]. Even eschewing an expressivist interpretation, the formal features and predictive power of these systems are of interest.

Our approach has precursors. Like us, Yalcin [2011] incorporates topic-sensitivity and fragmentation into his expressivism, with the explicit aim of evading pernicious transparency. We differ in key ways. First, Yalcin utilizes *domain semantics*.[11] While conditions for assertibility – 'acceptance', he calls it – are crucial for his account, they are defined derivatively. Second, Yalcin's system includes only belief ascriptions, leaving it (at best) unclear how its resources might assure apt properties for knowledge ascription. In this paper, knowledge ascription is the driving concern.

Yalcin [2012b] offers an alternative approach to evading pernicious transparency, rooted in a treatment of knowledge and belief states as *imprecise credal states*, i.e., sets of probability measures. However, this model shares key idealizations with Holism (e.g., true knowledge ascription is closed under entailment) that motivate, one might think, incorporation of topic-sensitivity and fragmentation. The moral of the present paper is that imprecise credence isn't *needed* for evading pernicious transparency:

[7] See Lewis [1988], Levesque and Lakemeyer [2000], Yalcin [2011, 2018], Yablo [2014], van Benthem [2011, Ch.6], Schipper [2015], Hawke [2016a,b, 2018], Berto [2018, 2019], Hawke et al. [2020], Berto and Hawke [2021], Plebani and Spolaore [2021], Hoek [2022, forthcoming].

[8] For an overview: Borgoni et al. [2021]. Also see Lewis [1982], Stalnaker [1984, Ch.4], Fagin and Halpern [1988], van Benthem [2011, Ch.5], Yalcin [2011], Rayo [2013, §4.3], Greco [2015], and Hawke et al. [2020].

[9] See Veltman [1985], Schroeder [2008b,a], Cantwell [2014, 2021], Hawke and Steinert-Threlkeld [2015, 2018, 2021], Lin [2016], Ciardelli [2021], [Aloni et al., 2023, Sect 6.1].

[10] Cf. Veltman [1996], Punčochář [2015, 2016], Fine [2017], Aloni [2022].

[11] See Dorr and Hawthorne [2013], Schroeder [2015] and Hawke and Steinert-Threlkeld [2021] for criticisms of domain semantics.

topic-sensitivity and fragmentation suffice, with a dose of contextualism.

Beddor and Goldstein [2021] (following Moss [2013]) develop another variant of domain semantics that explicitly evades pernicious transparency: building on Holism for belief ascription, they define knowledge as a conjunction of truth, belief and safety. A fair discussion requires some details; we leave this for Appendix A.

We proceed as follows. §2 informally assesses the challenge to expressivism from pernicious transparency and elaborates our proposed solution. §3 describes a basic holist assertibility semantics. §4 sharpens the challenge from pernicious transparency, by observing *prima facie* advantages of expressivist and descriptivist extensions of our basic semantics. §5 develops a fragmentationist and topic-sensitive expressivism that demonstrably delivers most such advantages, with general factivity one of two intriguing omissions. §6.1 motivates *hesitation* in accepting context-invariant factivity for modal knowledge ascriptions. §6.2 offers a contextualist refinement of our system that (i) rejects modal factivity in full generality, but with prospects for explaining away the intuitive judgments that supposedly support it, and (ii) evades a subtler form of problematic transparency. Appendix A evaluates Beddor and Goldstein's safety semantics; Appendix B considers how our system can model *modal Gettier cases* identified by Moss [2013].

2 Informal Backdrop and Basic Proposal

In this section, we elaborate key conceptual and motivational points informally. §2.1 clarifies the paper's use of 'expressivism'. §2.2 uses the challenge from pernicious transparency to frame a *would-be objection to expressivism*: accepting a pair of otherwise enticing commitments delivers pernicious transparency for the expressivist, saddling her with a counter-intuitive epistemic logic. The first commitment concerns the meaning of attitude ascriptions that embed an epistemic modal; the second concerns the explication of a key notion: *incompatibility with an acceptance state*. §2.3 motivates the paper's reply: resist the second commitment, by exploiting the topic-sensitivity and fragmentation of acceptance states. §2.4 sketches our leading ideas for an expressivism that promises a more sensible epistemic logic.

2.1 *Descriptivism versus expressivism*

Uttering the description 'it is raining' canonically communicates two bits of information. First, the representational content *that it is raining*; second, that the speaker is committed to the world being as that content represents. In general, we assume uttering description p in context c communicates both a canonical *proposition* (denoted $[p]_c$) and a canonical *state of mind*: commitment to $[p]_c$. We say that p *reports* $[p]_c$ and *expresses* commitment to $[p]_c$. The reporting part misleads if $[p]_c$ is false. The expression part misleads if the speaker isn't committed to $[p]_c$.

Note we *stipulate* 'proposition' to mean 'representational content'. Thus, we take propositions to have *subject matter* and to determine a *set of possible worlds*: those

that are as the proposition represents.[12] Hence, to communicate a proposition is to deliver *information* that excludes some ways the world could be. A proposition is *veridical* at world w iff its set of worlds includes w; two propositions are *consistent* iff there is a world at which they are both veridical; one proposition *entails* another iff every world where the first is veridical is one where the second is veridical.

What is *commitment* to a proposition? Consider the class of cognitive attitude verbs, including 'knows', 'believes', 'thinks', 'supposes', and 'assumes' (we take 'accepts' as generic). We take it as characteristic of every such attitude Φ that an agent's psychological state at time t includes a representation of the world (a set of propositions) aptly called her Φ-*state*. If proposition P is included in agent Smith's Φ-state, we say that Smith Φs P and that Smith's Φ-state is *committed* to P. For description p and context c, we assume that 'Smith Φs p' holds in c iff Smith Φs $[p]_c$. For example, 'Smith knows p' holds in c iff Smith knows $[p]_c$.

It eases our explication of *expression* to assume throughout that knowledge is both a mental state and the norm of assertion.

> **Knowledge Norm of Assertion (KNoA):** Declarative φ is assertible by agent a, in discourse context c, iff 'a knows φ' holds in c.

By φ being *assertible* by a, we mean that a is positioned to aptly assert φ in virtue of φ's literal meaning, the norms of assertion, and a's cognitive state (putting aside etiquette and Gricean pragmatics). Following Schroeder [2008b], we take 'expression' to track assertibility conditions: φ *expresses* state of mind Π exactly when Π is required for φ to be assertible. Given KNoA, this comes to: asserting φ in context c expresses that, in c, the speaker knows φ.

Given KNoA, descriptivists and expressivists agree: asserting might-p expresses that the speaker knows might-p. But they disagree on what knowing might-p comes to.[13]

> **Descriptivism For Bare Might Claims:** For every description p and context c, there exists proposition $[\text{might-}p]_c$ such that: in c, 'a knows might-p' holds iff a knows $[\text{might-}p]_c$.

> **Expressivism For Bare Might Claims:** For every description p and context c, there exists property Π_c such that: (i) Π isn't identical to merely knowing a proposition; (ii) in c, 'a knows might-p' holds iff a's knowledge state has Π_c.

[12] Presumably, this is the sense of 'proposition' that Dorr and Hawthorne [2013] and Beddor and Goldstein [2021] have in mind when labeling expressivism about might-p as *non-propositionalism*. For contexts where 'proposition' is intended to track pre-theoretic (meta-)semantic claims, some expressivists advocate pulling apart 'proposition' and 'representational content' (see Schroeder [2011], Ridge [2014], Bar-On et al. [2014]), citing minimalism or deflationism about propositions, or taking the proposition associated with an interpreted declarative as nearby its compositional semantic value, whatever this comes to. The purported advantage: expressivists can then claim that non-descriptive declaratives have *propositional* content, despite lacking *representational* content. Our discussion is consistent with this: such expressivists should understand our use of 'proposition' as technical, substituting 'representational content' if preferred.

[13] Prominent descriptivists include: Kratzer [1981, 2012], Dowell [2011], Silk [2017].

In cashing out this contrast in terms of *characteristic mental states*, we follow Gibbard [2003], Schroeder [2008a, Ch.1], and Yalcin [2011].[14] Going forward, 'descriptivism' is shorthand for the first thesis above; 'expressivism' for the second.

2.2 *The objection from transparency*

A straightforward and influential exemplar of expressivism says: for every description p and context c, 'a knows might-p' holds in c iff p is *compatible* with what a knows in c (compare a straightforward *descriptivism*: a knows might-p iff a knows that p is compatible with what a knows). What about cognitive attitude ascriptions *in general*? A natural expressivist move simply generalizes, following Yalcin [2007]:

> **Attitude Shift:** For any agent a and description p, 'a accepts that it might be that p' is equivalent to 'p is compatible with what a accepts'.

Again, 'accepts' is here *generic*: we take principles framed for acceptance as applying to *any* cognitive attitude verb. Note also that 'compatible' is here a technical (though suggestive) term. Here are two natural candidate definitions for 'p is compatible with what agent a accepts':

> **Com1:** 'compatibility' $=_{df}$ p is consistent with the sum total of what a accepts[15]

> **Com2:** 'compatibility' $=_{df}$ a doesn't accept not-p

Call an expressivist theory *Crude Expressivism* if it, like Yalcin [2007], either accepts Attitude Shift interpreted with Com1, or with Com2. Crude Expressivism entails:

> **Pernicious Transparency:** Either 'a accepts that it might be that p' is equivalent to 'p is consistent with the sum total of what a accepts', or to 'a doesn't accept not-p'.

But as §1 intimated, an objection looms: Crude Expressivism delivers a highly unorthodox epistemic logic, as noted by Yalcin [2012b, Sect.3], Dorr and Hawthorne [2013, fn.7], Moss [2018, Ch.6], and Beddor and Goldstein [2021, Sect. II.2]. Start with two orthodox principles of epistemic logic, where φ ranges over all declaratives:

> **Factivity:** 'a knows φ' entails 'φ'.[16]

> **KB:** 'a knows φ' entails 'a believes φ'.

For description p, we have the following special cases:

> **Modal Factivity:** 'a knows might-p' entails 'might-p'.

[14] Expressivism is compatible with some sentences that *embed* a bare might claim (e.g., not-might-p) being characteristically descriptive. It does not insist that 'might' is essentially a force-indicator. It is compatible with a *hybrid* view where might-p isn't a *mere* description, but loaded with further characteristic functions.

[15] This is plausibly the *default* notion of compatibility in the literature on modals/conditionals: compare [Kratzer, 2012, pg.11] and [Veltman, 1996, Sect. 2].

[16] Linguists call this 'veridicality': see Egré [2008].

Modal KB: 'a knows might-p' entails 'a believes might-p'.

If Pernicious Transparency holds, Modal Factivity and Modal KB are respectively equivalent to (6) and (7), or to (8) and (9):

(6) 'p is consistent with the sum total of what Smith knows' entails 'it might be that p'.

(7) 'p is consistent with the sum total of what Smith knows' entails 'p is consistent with the sum total of what Smith believes'.

(8) 'Smith doesn't know not-p' entails 'it might be that p'.

(9) 'Smith doesn't know not-p' entails 'Smith doesn't believe not-p'.

But all hands agree (6)-(9) are false. Suppose Smith is ignorant about the weather: for all she knows, it is raining, and for all she knows, it isn't. Knowing this about Smith doesn't typically permit one to conclude 'it might be raining': it isn't contradictory to deny that it is or might be raining while professing Smith's ignorance on the matter. Nor does Smith's ignorance entail that what Smith *believes* is consistent with it raining: she might unreasonably but firmly believe that it *isn't* raining.

So, Pernicious Transparency seemingly implies that Factivity and KB are false.

Further, any sensible account of knowledge ascription should reject the following (where as usual p ranges over descriptions):

Modal Omniscience: 'p' entails 'a knows might-p'.

Plutarch is a Greek philosopher, but Smith needn't even know that Plutarch *might* be a philosopher. Smith might falsely (but firmly and justifiably) believe that Plutarch isn't a philosopher, or might never have heard of Plutarch! But given Pernicious Transparency, Modal Omniscience is equivalent to one of:

(10) 'p' entails 'p is consistent with the sum total of what Smith knows'.

(11) 'p' entails 'Smith doesn't know not-p'.

All hands agree (10) and (11) are *true*, since false descriptions cannot be known. So, Pernicious Transparency egregiously implies Modal Omniscience is true.

In total, Crude Expressivism yields a highly unorthodox epistemic logic. This is untenable. But does the objection effectively scale up to attack expressivism *in general*? To survey strategies for resistance, consider this would-be objection:

The objection from transparency

P1. Expressivists must endorse Attitude Shift.

P2. Attitude Shift is best understood, via Com1/Com2, as Pernicious Transparency (thus, it entails Factivity and KB are false, and Modal Omniscience is true).

P3. At least one of the following holds: (i) Factivity is true, (ii) KB is true, or (iii) Modal Omniscience is false.

C. Thus, expressivism is false.

Is this cogent? Crude expressivists must, without much hope, deny P3: even Modal Factivity, the most questionable of (i)–(iii) (cf. §6.1), resists decisive rejection. For pre-theoretic evidence *against* Modal Factivity, an expressivist might exhibit:

(12) Smith knows that it might be raining and that it might not be raining.

(13) It might be raining and it might not be raining.

(12) predominantly communicates *Smith's* ignorance; (13) predominantly communicates *the speaker's* ignorance. If they aren't identical, Smith's ignorance doesn't in general imply anything about the speaker's ignorance. So couldn't (12) hold without (13), contravening Modal Factivity? The descriptivist has a ready reply: typical utterances of (12) and (13) illustrate the *context-sensitivity* of epistemic modals. Taking (12) and (13) as a counterexample to Factivity is akin to erroneously taking as a counterexample a true instance of 'Smith knows it is raining', uttered on a rainy day in Paris, coupled with a false instance of 'it is raining', uttered on a clear day in Cairo. The superiority of this explanation, the descripitivist adds, is witnessed by how jarring *bare* violations of Modal Factivity can sound:

(14) # Joe can't be the winner but Smith knows that Joe might be.

(15) Smith falsely believes Joe might win. # Indeed, she falsely knows he might.

Similarly, apparent violations of Modal KB sound jarring:

(16) # We all know Joe might win, but I doubt he might win.

(17) # Smith doesn't believe Joe might win, but she knows he might.

To my ears, (16)–(17) are most easily interpreted as reporting *cognitive dissonance*, in line with Modal KB (compare: 'Smith persists in believing Joe will win, but knows deep down he won't'). Compare benign claims: 'we all know Joe might win, but I doubt his chances'; 'Smith doesn't believe Joe is likely to win, but knows he might'.

Can sophisticated expressivists deny P1, denying Attitude Shift? Beddor and Goldstein [2021] follow this route, by treating knowledge as a composite of belief, safety, and truth. Appendix A explores this option. Another option modifies Attitude Shift along *hybrid expressivist* lines [Chrisman, 2007, Schroeder, 2009, Ridge, 2014], with the aim of *merging* features of basic descriptivism and basic expressivism. However, it is unclear what this should come to in the present setting, and whether it suffices to rule out a variant of the objection from transparency. Consider a simple hybrid variant of Attitude Shift: 'a accepts that it might be that p' is equivalent to the *conjunction* of 'a doesn't accept p' and 'a accepts that a doesn't know p'. Does this secure Modal KB? This requires that 'a doesn't know p' and 'a knows a doesn't know p' together entail both 'a doesn't believe p' and 'a believes a doesn't know p'. So, it requires: 'a knows a doesn't know p' entails 'a doesn't believe p' (assuming 'a knows a doesn't know p' entails 'a doesn't know p'). But this is false: one can know that one's evidence doesn't position one to *know* that is raining in London, yet

one nevertheless believes (tentatively or irrationally) that it is raining in London.

The current paper thus explores the remaining strategy for resisting the objection from transparency: deny P2, by resisting definitions Com1 or Com2 of 'compatibility'.

2.3 *Holism, topic-sensitivity, and fragmentation*

Compare further candidate definitions for 'p is compatible with what agent a accepts':

> **Com3:** 'compatibility' $=_{df}$ a doesn't accept not-p despite a's acceptance state including content that is partly about p's subject matter (we say a's acceptance state is *sensitive* to p's subject matter)
>
> **Com4:** 'compatibility' $=_{df}$ a doesn't accept not-p in at least one frame of mind
>
> **Com5:** 'compatibility' $=_{df}$ a doesn't accept not-p in at least one frame of mind that is partly about p's subject matter

Com1-Com5 naturally partner with influential approaches to modeling acceptance states and acceptance attributions.

- **Holism:**[17] An agent's acceptance state is best modeled as a single proposition; a proposition is best modeled as a set of possible worlds; 'a accepts p' holds in c exactly when a's acceptance state entails $[p]_c$.
- **Topic-sensitivity:**[18] A proposition is best modeled as a set of possible worlds (representing verification/truth conditions) plus a subject matter (a topic or question, or set of such things, that the proposition *addresses* or is *about*); 'a accepts p' holds only if the content of a's acceptance state is partly about p's subject matter.
- **Fragmentation:**[19] An agent's acceptance state is best modeled as a set of propositions called *frames of mind*; 'a accepts p' holds just in case p is supported by at least one of a's frames of mind.

Com1 and Com2 are equivalent to a natural Holist view: p is compatible with an acceptance state **s** iff p holds at a world in **s**. Com3 complements Topic-sensitivity. Com4 complements Fragmentation. Com5 fits the natural *combination* of Topic-sensitivity and Fragmentation. So, choosing a definition relates intimately to choosing a model of acceptance states.

Holism's appeal lies in its mathematical elegance and prediction that attitude ascriptions cohere in intuitive ways: 'Smith knows there is a black sheep' entails 'Smith

[17] Cf. Hintikka [1962].

[18] Cf. Lewis [1988], Fagin and Halpern [1988], Yablo [2014, Ch.7], Hawke [2016a,b], Berto [2019], Hawke et al. [2020].

[19] Cf. Lewis [1982], Stalnaker [1984, Ch.4], Fagin and Halpern [1988], Yalcin [2011, 2018], Hawke et al. [2020].

knows there is a sheep', and so on. However, it also notoriously entails logical omniscience and the impossibility of cognitive dissonance. Topic-sensitivity and Fragmentation preserve Holism's virtues while alleviating its vices, with refinements inspired by plausible features of ordinary cognition. As propositional content has subject matter (determining the topics it is about), ordinary acceptance states are plausibly topic-sensitive. Ann might accept that Bob isn't a lawyer ($\neg p$), without accepting something logically equivalent: that Bob isn't a lawyer *and* Bob isn't both a lawyer and an admirer of Hausdorff's contributions to topology ($\neg p \wedge \neg(p \wedge q)$). After all, if Ann has never have heard of Hausdorff, or lacks the concept 'topology', she can't even *entertain* $\neg p \wedge \neg(p \wedge q)$. Or perhaps Ann has heard of Hausdorff and topology, but it has never occurred to her to think about them and Bob in concert. Either way, Ann's acceptance state is 'insensitive' to the subject matter of $\neg p \wedge \neg(p \wedge q)$: she neither ponders, accepts, rejects, nor explicitly suspends judgment on $\neg p \wedge \neg(p \wedge q)$.

Ordinary acceptance states are also plausibly fragmented. Adam might know a rectangular plot of land is 16m long and 7m wide without knowing its area is $112\mathrm{m}^2$, despite grasping the concepts needed for the calculation. Intuitively, he is yet to properly *integrate* his knowledge of the plot's dimensions with his arithmetical knowledge (cf. [Rayo, 2013, Sect. 4.3]). Perhaps his attempted calculation went astray: he believes, inconsistently, that the area is $110\mathrm{m}^2$. Or consider Eve: she simultaneously believes and disbelieves a *single* content, indexed to different presentations: 'lung cancer patients with a 90% one-month survival rate should get surgery' elicits her endorsement while 'lung cancer patients should get surgery with a 10% one-month mortality rate' elicits her opposition [Kahneman and Tversky, 1984, Berto and Özgün, 2023]. Adam and Eve are readily understood as having frames of mind – belief fragments – with content that is individually consistent, but jointly inconsistent.

Those embracing Topic-sensitivity or Fragmentation needn't commit to Com1 and Com2 – to their advantage. Consider:

(18) Smith believes might-p.

(19) Smith doesn't believe not-p.

(20) Smith doesn't believe not-p yet her belief state is sensitive to p's subject matter.

(21) In *one* frame of mind, Smith doesn't believe not-p.

Given Com2, Attitude Shift entails (18) and (19) are equivalent; given Com3, it entails (18) and (20) are equivalent. Yalcin [2011] notes Topic-sensitivity has here a pre-theoretic edge over Holism: intuitively, if Smith hasn't heard of Topeka, 'Smith doesn't believe that it isn't raining in Topeka' holds without 'Smith believes that it might be raining in Topeka'. Given Com4, Attitude Shift entails (18) and (21) are equivalent. Fragmentation has here a pre-theoretic edge over Holism: just as cognitive dissonance allows 'Smith believes all life is sacred' and 'Smith believes a slug's life is worthless' to hold simultaneously, it presumably allows 'Smith believes all life is sacred' (entailing 'Smith believes a slug's life isn't worthless') and 'Smith

believes a slug's life *might* be worthless' to hold simultaneously, across different frames of mind.[20] So, proponents of Attitude Shift have *prima facie* reason to adopt *both* Topic-Sensitivity and Fragmentation, combining the virtues of Com3 and Com4 in Com5.

2.4 *Our proposed framework in outline*

Our hypothesis: proponents of Attitude Shift can dodge (the egregious consequences of) Pernicious Transparency by refining Holism with Fragmentation and Topic-sensitivity. Coming sections develop this into a rigorous theory. Here are the key ideas.

We model knowledge and belief as fragmented and assume that a knowledge fragment is a *type* of belief fragment (remaining silent on the epistemological question as to what *makes* a belief fragment a knowledge fragment). Attitude attribution follows the fragmentationist template (cf. [Fagin and Halpern, 1988]): 'Smith knows that φ' holds when at least one knowledge fragment supports φ; ditto for belief attribution. This ensures KB: if a knowledge fragment supports φ, so does a belief fragment.

We model propositional content with two components (cf. [Yalcin, 2011, Yablo, 2014, Berto, 2019]): verification conditions and subject matter. Thus, knowledge and belief are topic-sensitive: for certain subject matters, the agent's state contains fragments that speak to that subject matter; for others, there may be no such fragment. Attitude attribution follows the topic-sensitive template: 'Smith knows that φ' holds only if Smith's knowledge state includes content about whatever φ is about. This sabotages Modal Omniscience: just because φ holds doesn't guarantee that Smith's knowledge state is sensitive to its subject matter.

This theory rejects (key consequences of) Pernicious Transparency. But we will confirm a loose end: topic-sensitivity and fragmentation do *not* by themselves guarantee Modal Factivity. The tension between expressivism and factivity runs deep. Fortunately for expressivists, §6.1 showcases pre-theoretic evidence against Modal Factivity. §6.2 shows that expressivists can account for *both* this data *and* data (e.g., (14), (15)) that conforms to Modal Factivity, by predicting that Modal Factivity holds in a prominent but *restricted* set of contexts.

3 Holistic Assertibility Semantics

We shift to a technical mode, working in the setting of *assertibility semantics*.[21] We start with a basic holist semantics. We modify it in coming sections, to frame both the challenge from pernicious transparency and our proposed solution more precisely.

Concentrating on assertibility has an important advantage in the present debate: relative *neutrality*. Disagreement is wide on how to represent the fundamental semantic

[20] MacFarlane [2014, pg. 279] labels this 'epistemic *akrasia*'.

[21] Expressivist approaches: Veltman [1985], Schroeder [2008b], Lin [2016], Hawke and Steinert-Threlkeld [2018], Cantwell [2014], Hawke and Steinert-Threlkeld [2021], Ciardelli [2021]. Adjacent work: Veltman [1996], Punčochář [2015, 2016], Aloni [2022], [Aloni et al., 2023, Sect 6.1].

values of declaratives. Descriptivists favour a traditional truth-conditional semantics, broadly construed [Kratzer, 1981, Silk, 2017]. Expressivism fits better with domain semantics [Yalcin, 2007], dynamic semantics [Yalcin, 2012b, Willer, 2013, 2017b], inferentialist semantics [Incurvati and Schlöder, 2019], or a 'psychologistic semantics' where assertiblity conditions are fundamental [Schroeder, 2008b, Hawke and Steinert-Threlkeld, 2021, Ciardelli, 2021]. But all hands acknowledge that assertibility conditions are usefully studied, even if derivative of fundamental semantic facts.

3.1 *Basics for assertibility semantics*

A formal assertibility semantics models the *assertibility relation* $\Vdash$, holding between a unified body of information **s** (an *information state*) and a meaningful declarative φ, exactly when: were an agent's knowledge state to contain exactly information **s**, she would be correct to assert φ, from a purely semantic and epistemic perspective (cf. [Jackson, 1987, pg.8], [Gauker, 2005, Sect. 9.2], [Schroeder, 2008b]). We assume that an information state can be identified with a proposition and use $\mathbf{I}$ to denote the set of all information states. Call a subset of $\mathbf{I}$ a *cognitive feature*.

We model declaratives with formal language $\mathcal{L}$, including countable *atomic descriptions*; negation $\neg$; conjunction $\wedge$; 'might' operator $\diamond$; belief operator B; knowledge operator K. Read $\diamond\varphi$ as 'it might be that φ'; $B\varphi$ as 'Smith believes φ'; $K\varphi$ as 'Smith knows φ'. We use φ, ψ as meta-variables over $\mathcal{L}$-sentences, and p, q, r over atoms.

We interpret $\mathcal{L}$ on different flavors of *epistemic frame*, a tuple with at least: a set of all possible worlds W; a designated actual world @; a belief function; a knowledge function. The latter two map each possible world to an *acceptance state*: respectively, Smith's *belief state* at w and *knowledge state* at w. We call a subset of W an *intension*. A *valuation function* v assigns a proposition $\mathrm{v}(p)$ to each atom p. An epistemic frame with a valuation function is an *epistemic model*, denoted $\mathcal{M}$.

We'll see competing proposals for precisifying 'epistemic model' and defining $\Vdash$. Call this combination a *system (for $\mathcal{L}$)* with an *account* of $\Vdash$.

Definition 1 (**Expressed Feature**). *Relative to a model and an account of $\Vdash$, the cognitive feature expressed by φ is:* $[\![\varphi]\!] := \{\boldsymbol{s} \in \mathbf{I} : \boldsymbol{s} \Vdash \varphi\}$.

So, φ expresses the type of information state that renders φ assertible.

3.2 *Holist semantics*

How to model propositions, information states, and acceptance states? Holism says: identify each with an intension, thereby precisifying 'epistemic frame' and 'epistemic model' as *holist frames* and *holist models* (denoted $\mathcal{H}$). The intension corresponding to a proposition is intuitively the set of worlds left uneliminated if that proposition is accepted. Thus, in a holist frame, belief function **b** and knowledge function **k** map each world to an intension. We stipulate: for every w in a holist model, $\mathbf{b}(w) \subseteq \mathbf{k}(w)$ (an agent's belief state eliminates at least as much as her knowledge state) and $w \in \mathbf{k}(w)$ (knowledge states can't rule out 'actuality').

Now for a holist account of $\Vdash$. The definition is *bilateral*: we mutually recursively define $\Vdash$ and the *deniability relation* $\dashv\Vdash$. Read '$\mathbf{s} \Vdash \varphi$' as '$\varphi$ is assertible given $\mathbf{s}$' or '$\mathbf{s}$ accepts φ'; ditto for $\dashv\Vdash$, 'deniable', and 'rejects'. Note $\mathbf{s} \nVdash \varphi$ means: 'not $\mathbf{s} \Vdash \varphi$'.

Definition 2 (Holism). *For $\boldsymbol{s}$, p, φ and ψ, relative to holist model $\mathcal{H}$:*

$\boldsymbol{s} \Vdash p$	*iff*	$\boldsymbol{s} \subseteq \mathrm{v}(p)$
$\boldsymbol{s} \dashv\Vdash p$	*iff*	$\boldsymbol{s} \cap \mathrm{v}(p) = \varnothing$
$\boldsymbol{s} \Vdash \neg\varphi$	*iff*	$\boldsymbol{s} \dashv\Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash \neg\varphi$	*iff*	$\boldsymbol{s} \Vdash \varphi$
$\boldsymbol{s} \Vdash \varphi \wedge \psi$	*iff*	$\boldsymbol{s} \Vdash \varphi$ *and* $\boldsymbol{s} \Vdash \psi$
$\boldsymbol{s} \dashv\Vdash \varphi \wedge \psi$	*iff*	*there are* $\boldsymbol{u}$ *and* $\boldsymbol{v}$ *s.t.* $\boldsymbol{s} = \boldsymbol{u} \cup \boldsymbol{v}$ *and* $\boldsymbol{u} \dashv\Vdash \varphi$ *and* $\boldsymbol{v} \dashv\Vdash \psi$

Description p is assertible exactly when the given information leaves only p-worlds uneliminated; $\neg p$ is assertible exactly when the given information eliminates all p-worlds; $\neg\varphi$ is assertible exactly when φ is deniable, and *vice versa*; $\varphi \wedge \psi$ is assertible exactly when both conjuncts are assertible; $\neg(\varphi \wedge \psi)$ is assertible exactly when the given information can be divided into two *jointly exhaustive* refinements, one rejecting φ, one rejecting ψ (so, $\neg(p \wedge q)$ is assertible exactly when any world uneliminated by the given information is either a $\neg p$-world or a $\neg q$-world).

The above clauses essentially match propositional team logic [Yang and Väänänen, 2017].[22] Notation aside, this foundation has been deployed to study of modals and conditionals by Lin [2016], Hawke and Steinert-Threlkeld [2015, 2018], Hawke and Steinert-Threlkeld [2021, Sect.5.1], Ciardelli [2021, Sects. 3, 5.2], and Aloni [2022, Sect. 4].[23] Related accounts include [Gauker, 2005, Ch.5] and [Cantwell, 2014, 2021].

Definition 3 (Holism continued).

$\boldsymbol{s} \Vdash K\varphi$	*iff*	$\forall w \in \boldsymbol{s}$: $\boldsymbol{k}(w) \Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash K\varphi$	*iff*	$\forall w \in \boldsymbol{s}$: $\boldsymbol{k}(w) \nVdash \varphi$
$\boldsymbol{s} \Vdash B\varphi$	*iff*	$\forall w \in \boldsymbol{s}$: $\boldsymbol{b}(w) \Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash B\varphi$	*iff*	$\forall w \in \boldsymbol{s}$: $\boldsymbol{b}(w) \nVdash \varphi$

By holist lights, $K\varphi$ and $B\varphi$ are thus descriptions: 'Smith knows that p' is assertible by a speaker with knowledge state $\mathbf{s}$ exactly when Smith's knowledge state $\mathbf{k}(w)$ renders p assertible at every world w uneliminated by $\mathbf{s}$, i.e., the speaker knows proposition $\{w \in W : \mathbf{k}(w) \Vdash \varphi\}$. (We do *not* assume Smith is always the speaker.)

By duality, disjunction is *split disjunction*:[24]

$\mathbf{s} \Vdash \varphi \vee \psi$	iff	there are $\mathbf{u}$ and $\mathbf{v}$ s.t. $\mathbf{s} = \mathbf{u} \cup \mathbf{v}$ and $\mathbf{u} \Vdash \varphi$ and $\mathbf{v} \Vdash \psi$
$\mathbf{s} \dashv\Vdash \varphi \vee \psi$	iff	$\mathbf{s} \dashv\Vdash \varphi$ and $\mathbf{s} \dashv\Vdash \psi$

[22] The system in [Yang and Väänänen, 2017] is unilateral; for basic propositional logic without expressivistic clauses for epistemic modals, this is equivalent to our bilateral system.

[23] The system in Ciardelli [2021] does not include clauses for conjunction.

[24] Cf. Lin [2016], Yang and Väänänen [2017], Hawke and Steinert-Threlkeld [2015, 2021], Aloni [2022].

3.3 *Assertoric logic*

With a system on the table (holist or otherwise), logical notions can be defined. As we assume the knowledge norm of assertion, we focus on the logic of assertibility at *veridical* (i.e. knowledge-like) states (for holists, $@ \in \mathbf{s}$).

Definition 4 (**Assertoric Consequence**). *$\varphi \Vdash \psi$ holds iff, for every veridical $\boldsymbol{s}$ in every $\mathcal{M}$, if $\boldsymbol{s} \Vdash \varphi$ then $\boldsymbol{s} \Vdash \psi$.*

So, $\Vdash$ indicates assertibility preservation across any conceivable epistemic scenario. We say that φ *(assertorically) entails* ψ when $\varphi \Vdash \psi$; *(assertoric) equivalence* is two-way entailment: $\varphi \dashv\Vdash \psi$. For example, our holist semantics agreeably validates:

Descriptive KB:[25] $Kp \Vdash Bp$

Descriptive Factivity:[26] $Kp \Vdash p$

We say: φ and ψ are *co-assertible* when there exists $\mathcal{M}$ and veridical $\mathbf{s}$ such that $\mathbf{s} \Vdash \varphi$ and $\mathbf{s} \Vdash \psi$; φ and ψ are *(assertorically) contrary* when not co-assertible; φ is an *(assertoric) contradiction* (written $\dashv\Vdash \varphi$) when there is no $\mathcal{M}$ and veridical $\mathbf{s}$ for which $\mathbf{s} \Vdash \varphi$. If necessary, we write $\underset{S}{\Vdash}$ and $\underset{S}{\dashv\Vdash}$ to highlight that system S's account of $\Vdash$ and $\dashv\Vdash$ is at issue.

4 Sharpening the Challenge from Transparency

We now compare the advantages of expressivist and descriptivist extensions of our holist semantics, resulting in a sharpened challenge from pernicious transparency.

4.1 *Holistic expressivism*

Consider a quintessentially expressivist extension of our holist semantics:

Definition 5 (**Holistic Expressivism (HEx)**).

$$\textbf{\textit{Holism}} + \quad \begin{array}{lll} \boldsymbol{s} \Vdash \Diamond\varphi & \textit{iff} & \boldsymbol{s} \not\dashv\Vdash \varphi \\ \boldsymbol{s} \dashv\Vdash \Diamond\varphi & \textit{iff} & \boldsymbol{s} \dashv\Vdash \varphi \end{array}$$

Hence, $\Diamond p$ is assertible given information $\mathbf{s}$ exactly when p is consistent with $\mathbf{s}$, i.e., there is a p-world in $\mathbf{s}$ (cf. Com1, Com2). Given the knowledge norm of assertion, we thus have a form of standard expressivism: a speaker knows $\Diamond p$ exactly when $\Diamond p$ is assertible given her knowledge state, which is exactly when p is consistent with that knowledge state. Notation aside, the above clauses are utilized in [Hawke and Steinert-Threlkeld, 2021, Sect. 5.1] and [Ciardelli, 2021, Sect. 3], and are in the spirit of Incurvati and Schlöder [2019], where uttering $\Diamond p$ is taken to communicate the 'weak assertion' of p (i.e., that p can't be rejected).

[25] *Proof.* Given **Holism**, assume $\mathbf{s} \Vdash Kp$. So, $\forall w \in \mathbf{s}$: every world in $\mathbf{k}(w)$ is a p-world. But per **Holism**: $\mathbf{b}(w) \subseteq \mathbf{k}(w)$. So, $\forall w \in \mathbf{s}$: every world in $\mathbf{b}(w)$ is a p-world. So, $\mathbf{s} \Vdash Bp$.

[26] *Proof.* Given **Holism**, assume $\mathbf{s} \Vdash Kp$. So, $\forall w \in \mathbf{s}$: every world in $\mathbf{k}(w)$ is a p-world. But per **Holism**: $w \in \mathbf{k}(w)$. So, every world in $\mathbf{s}$ is a p-world. So, $\mathbf{s} \Vdash p$.

The resulting notions of entailment and contradiction yield familiar expressivist advantages, transferring directly to HEx's logic of attitude ascriptions.

Proposition 1 (Inheritance). *According to HEx, $\mathbf{s} \Vdash p \wedge \diamond q$ implies $\mathbf{s} \Vdash \diamond(p \wedge q)$, for every holist model $\mathcal{H}$ and information state $\mathbf{s}$. In particular: $p \wedge \diamond q \underset{\texttt{HEx}}{\Vdash} \diamond(p \wedge q)$.*

Proof. Assume $\mathbf{s} \Vdash p \wedge \diamond q$. Then, $\mathbf{s} \Vdash p$ and $\mathbf{s} \not\Vdash q$. Thus, $\mathbf{s} \subseteq \mathrm{v}(p)$ and $\mathbf{s} \cap \mathrm{v}(q) \neq \varnothing$. So, $\exists w \in \mathbf{s}$ s.t. $w \in \mathrm{v}(p) \cap \mathrm{v}(q)$. So, $\mathbf{s} \not\Vdash (p \wedge q)$. So, $\mathbf{s} \Vdash \diamond(p \wedge q)$. □

This predicts the intuitive entailment of (23) from (22) [Dorr and Hawthorne, 2013]:

(22) I'm staying home and might watch a movie.

(23) I might stay home and watch a movie.

Importantly, the 'inheritance' phenomenon survives embeddings.

Proposition 2. $p \vee (q \wedge \diamond r) \underset{\texttt{HEx}}{\Vdash} p \vee \diamond(q \wedge r)$

Proof. Assume $\mathbf{s} \Vdash p \vee (q \wedge \diamond r)$. So, there exists $\mathbf{u}$ and $\mathbf{v}$ s.t. $\mathbf{s} = \mathbf{u} \cup \mathbf{v}$ and $\mathbf{u} \Vdash p$ and $\mathbf{v} \Vdash q \wedge \diamond r$. By our previous result: $\mathbf{u} \Vdash p$ and $\mathbf{v} \Vdash \diamond(q \wedge r)$. So, $\mathbf{s} \Vdash p \vee \diamond(q \wedge r)$. □

Again a pleasing prediction, as (24) intuitively implies (25):

(24) Tonight, I'm either bar-hopping or I'm staying home and might watch a movie.

(25) Tonight, I'm either bar-hopping or I might watch a movie at home.

As $\diamond$-claims are closed under entailment and inheritance survives substitution of complex descriptions, we also get:

$$\neg(p \wedge q) \wedge \diamond p \underset{\texttt{HEx}}{\Vdash} \diamond(\neg(p \wedge q) \wedge p) \underset{\texttt{HEx}}{\Vdash} \diamond\neg q$$

Indeed, the following reasoning is intuitively valid [Bledin and Lando, 2018, §3]:

(26) Ed didn't both study hard and fail. Ed might have studied hard. So, Ed might have passed.

Finally, inheritance renders $p \wedge \diamond\neg p$ an assertoric contradiction (that uniformly *embeds* like a contradiction), as from $p \wedge \diamond\neg p$ one concludes $\diamond(p \wedge \neg p)$. Indeed, the following has an intuitively contradictory air (cf. [Mandelkern, 2019]):

(27) # Either it's raining and might not be, or it's cloudy and might not be.

One can easily show that the above assertibility principles transfer directly to HEx's logic of attitude reports.

Corollary 1 (Attitude Inheritance).

$B(p \wedge \diamond q) \underset{\texttt{HEx}}{\Vdash} B\diamond(p \wedge q)$
$K(p \wedge \diamond q) \underset{\texttt{HEx}}{\Vdash} K\diamond(p \wedge q)$

Corollary 2 (Modal Syllogism).

$B(\neg(p \wedge q) \wedge \diamond p) \underset{\texttt{HEx}}{\Vdash} B \diamond \neg q$
$K(\neg(p \wedge q) \wedge \diamond p) \underset{\texttt{HEx}}{\Vdash} K \diamond \neg q$

Corollary 3 (Attitude Contradiction).

$\underset{\texttt{HEx}}{\nVdash} B(p \wedge \diamond \neg p)$
$\underset{\texttt{HEx}}{\nVdash} K(p \wedge \diamond \neg p)$

Thus, HEx matches intuition in predicting the incoherence of the following:

(28) # Smith knows Ed is old and might be grumpy, but not that Ed might be both old and grumpy.

(29) # Smith knows Ed might order champagne and never orders it without ordering orange juice, but doesn't know Ed might order orange juice.

(30) # Smith believes Ed is tall and might not be.

Turn to HEx's *prima facie* disadvantages regarding attitude reports. HEx directly validates symmetric transparency principles, exemplifying Crude Expressivism.

Proposition 3 (K-Transparency). $K \diamond \varphi \underset{\texttt{HEx}}{\Vdash\!\!\Vdash} \neg K \neg \varphi$

Proof. Assume $\mathbf{s} \Vdash K \diamond \varphi$. So, $\forall w \in \mathbf{s}$, $\mathbf{k}(w) \Vdash \diamond \varphi$ (i.e., $\mathbf{k}(w) \not\Vdash \varphi$). So, $\forall w \in \mathbf{s}$, $\mathbf{k}(w) \not\Vdash \neg \varphi$. So, $\mathbf{s} \Vdash \neg K \neg \varphi$. The reasoning can be reversed. □

Proposition 4 (B-Transparency). $B \diamond \varphi \underset{\texttt{HEx}}{\Vdash\!\!\Vdash} \neg B \neg \varphi$

Proof. Similar to **K-Transparency**. □

As critics forecast (cf. §2), these flag that HEx (disconcertingly) invalidates:

KB: $K\varphi \Vdash B\varphi$ **Factivity:** $K\varphi \Vdash \varphi$

To see this, focus on:

Modal KB: $K \diamond p \Vdash B \diamond p$ **Modal Factivity:** $K \diamond p \Vdash \diamond p$

Now note that HEx (matching intuition) delivers the following.

Proposition 5 (Contraposed BK fails). $\neg K \neg p \underset{\texttt{HEx}}{\not\Vdash} \neg B \neg p$

Proof. Counterexample: a holist model $\mathcal{H}$ and veridical $\mathbf{s}$ where $\mathbf{b}(@)$ only contains $\neg p$-worlds, but $\mathbf{k}(w)$ contains a p-world for every $w \in \mathbf{s}$. □

Proposition 6 (K$\diamond$ fails). $\neg K \neg p \underset{\texttt{HEx}}{\not\Vdash} \diamond p$

Proof. Counterexample: a holist model $\mathcal{H}$ and veridical **s** where **s** contains only $\neg p$-worlds, but $\mathbf{k}(w)$ contains a p-world for every $w \in \mathbf{s}$. □

But **Modal KB** plus **K-Transparency** and **B-Transparency** together entail **Contraposed BK**. Consider the chain of implications:

$$\neg K \neg p \Vdash K \diamond p \Vdash B \diamond p \Vdash \neg B \neg p$$

Hence, HEx doesn't validate **Modal KB**. Furthermore, **Modal Factivity** plus **K-Transparency** and **B-Transparency** entails **K**$\diamond$:

$$\neg K \neg p \Vdash K \diamond p \Vdash \diamond p$$

So, HEx doesn't validate **Modal Factivity**.

Furthermore, HEx problematically validates:

Modal Omniscience: $p \Vdash K \diamond p$

This again relates closely to **K-Transparency**, as HEx (agreeably) delivers:

Proposition 7 (**Contraposed Descriptive Factivity**). $p \underset{\texttt{HEx}}{\Vdash} \neg K \neg p$

Proof. Assume $\mathbf{s} \Vdash p$. So, $\forall w \in \mathbf{s}$: w is a p-world and so $\mathbf{k}(w) \nVdash \neg p$, as $w \in \mathbf{k}(w)$. So, $\forall w \in \mathbf{s}$: $\mathbf{k}(w) \Vdash \neg K \neg p$. So, $\mathbf{s} \Vdash \neg K \neg p$. □

Finally, **K-Transparency**, interacting with **Attitude Inheritance**, validates another undesirable (and closely related) principle:

Uncertainty: $K(p \vee q) \wedge \neg K p \Vdash K \diamond q$

To see this, note that HEx closes K and $\diamond$ under entailment and also validates:

Agglomeration: $K\varphi \wedge K\psi \Vdash K(\varphi \wedge \psi)$

Now consider the chain of implications:

$$K(p \vee q) \wedge \neg K p \Vdash K(p \vee q) \wedge K \diamond \neg p \Vdash K \diamond ((p \vee q) \wedge \neg p) \Vdash K \diamond q$$

Uncertainty has intuitive counterexamples. Suppose Smith knows a coin landed Heads or Tails, but not which. She won if it landed Tails (it did). But she mistakenly believes the coin is a trick coin that never lands Tails, so firmly *believes* that it landed Heads. So, she neither believes, nor knows, that it might have landed Tails.

4.2 *Holistic descriptivism*

Alternatively, one could extend our holist semantics as follows:

Definition 6 (**Holistic Descriptivism (HolD)**).

Holism +
$$\begin{array}{lll} \boldsymbol{s} \Vdash \diamond\varphi & \textit{iff} & \forall w \in \boldsymbol{s}\textit{:}\ \boldsymbol{k}(w) \not\dashv \varphi \\ \boldsymbol{s} \dashv \diamond\varphi & \textit{iff} & \forall w \in \boldsymbol{s}\textit{:}\ \boldsymbol{k}(w) \dashv \varphi \end{array}$$

Hence, $\diamond p$ is assertible given $\mathbf{s}$ when p is consistent with the relevant agent's (Smith's) knowledge state at every world uneliminated by $\mathbf{s}$. By holist lights, $\diamond p$ expresses a speaker's commitment to the proposition $\{w \in W : \mathbf{k}(w) \not\Vdash \varphi\}$, yielding a straightforward descriptivism (via Com1 or Com2). Restricted to description p, the above clauses formally match the Bilateral State-based Modal Logic of Aloni [2022, Sect.4].

With respect to attitude reports, HolD's strengths and weaknesses are the *mirror image* of HEx, at least for the properties in our focus.

Proposition 8 (Modal KB). $K \diamond p \underset{\texttt{HolD}}{\Vdash} B \diamond p$

Proof. Assume $\mathbf{s} \Vdash K \diamond p$. So, $\forall w \in \mathbf{s}$: $\mathbf{k}(w) \Vdash \diamond p$. So, $\forall w \in \mathbf{s}$: $\forall v \in \mathbf{k}(w)$: $\mathbf{k}(v) \not\Vdash p$. So, $\forall w \in \mathbf{s}$: $\forall v \in \mathbf{k}(w)$: $\mathbf{k}(v) \cap \mathrm{v}(p) \neq \varnothing$. Since $\mathbf{k}(v) \subseteq \mathbf{b}(v)$ for all v, it follows that: $\forall w \in \mathbf{s}$: $\forall v \in \mathbf{k}(w)$: $\mathbf{b}(v) \cap \mathrm{v}(p) \neq \varnothing$. So, $\forall w \in \mathbf{s}$: $\mathbf{b}(w) \Vdash \diamond p$. So, $\mathbf{s} \Vdash B \diamond p$. □

Proposition 9 (Modal Factivity). $K \diamond p \underset{\texttt{HolD}}{\Vdash} \diamond p$

Proof. Assume $\mathbf{s} \Vdash K \diamond p$. So, $\forall w \in \mathbf{s}$: $\forall v \in \mathbf{k}(w)$: $\mathbf{k}(v) \not\Vdash p$. It follows, as $w \in \mathbf{k}(w)$, that $\forall w \in \mathbf{s}$: $\mathbf{k}(w) \not\Vdash p$. So, $\mathbf{s} \Vdash \diamond p$. □

Indeed, HolD validates **KB** and **Factivity** in general. Notably, it does so while agreeably delivering:

Proposition 10 (Contraposed BK fails). $\neg K \neg p \underset{\texttt{HolD}}{\not\Vdash} \neg B \neg p$.

Proof. Counterexample: model $\mathcal{H}$ and veridical $\mathbf{s}$ where $\mathbf{b}(@)$ only contains $\neg p$-worlds, $\mathbf{k}(@)$ contains a p-world, and $\forall w \in \mathbf{s}$: $\mathbf{b}(w) = \mathbf{b}(@)$ and $\mathbf{k}(w) = \mathbf{k}(@)$. □

As **Modal KB** plus **K-Transparency** and **B-Transparency** entails **Contraposed BK**, it follows that HolD agreeably invalidates **K-transparency** and **B-Transparency**.

Similarly, HolD agreeably delivers:

Proposition 11 (Modal Omniscience fails). $p \underset{\texttt{HolD}}{\not\Vdash} K \diamond p$.

Proof. Counterexample: any model $\mathcal{H}$ (evaluated at $\mathbf{s} = \{@\}$) where $@$ is a p-world and $\exists w \in \mathbf{k}(@)$ s.t every world in $\mathbf{k}(w)$ is a $\neg p$-world. □

Notably, it does so while delivering:

Proposition 12 (Contraposed Descriptive Factivity). $p \underset{\texttt{HolD}}{\Vdash} \neg K \neg p$.

Proof. Assume $\mathbf{s} \Vdash p$. So, every $w \in \mathbf{s}$ is a p-world. So, $\forall w \in \mathbf{s}$: $\mathbf{k}(w)$ contains a p-world (namely, w). So, $\forall w \in \mathbf{s}$: $\mathbf{k}(w) \not\Vdash \neg p$. So, $\mathbf{s} \Vdash \neg K \neg p$. □

Finally:

Proposition 13 (Uncertainty fails). $K(p \vee q) \wedge \neg Kp \underset{\texttt{HolD}}{\not\Vdash} K \diamond q$.

Proof. Counterexample: $\mathcal{H}$ where: (i) $W = \{w_1, w_2\}$, (ii) $\mathrm{v}(p) = \{w_2\}$, (iii) $\mathrm{v}(q) = \{w_1\}$, (iv) $\mathbf{k}(w_1) = \{w_1, w_2\}$, (v) $\mathbf{k}(w_2) = \{w_2\}$, and (vii) $@ = w_1$. Let $\mathbf{s} = \{@\}$. □

However, on the negative side, HolD problematically delivers the following.

Proposition 14 (Attitude Contradiction fails). $\underset{\texttt{HolD}}{\nVdash} B(p \wedge \diamond \neg p)$.

Proof. Note: $\{@\} \Vdash B(p \wedge \diamond \neg p)$ iff $\mathbf{b}(@) \Vdash p \wedge \diamond \neg p$. So, one can emulate a proof that $\mathbf{s} \Vdash p \wedge \diamond \neg p$ can hold to conclude that $\{@\} \Vdash B(p \wedge \diamond \neg p)$ can hold. □

Proposition 15 (Modal Syllogism fails). $B(\neg(p \wedge q) \wedge \diamond p) \underset{\texttt{HolD}}{\not\Vdash} B \diamond \neg q$

Proof. Counterexample: $\mathcal{H}$ where: (i) $W = \{w_1, w_2\}$, (ii) $\mathrm{v}(p) = \{w_1\}$, (iii) $\mathrm{v}(q) = W$, (iv) $\mathbf{b}(w_1) = \mathbf{b}(w_2) = \{w_2\}$, and (v) $\mathbf{k}(w_1) = \mathbf{k}(w_2) = W$. Let $\mathbf{s} = W$. □

As a consequence, **Attitude Inheritance** fails: $B(p \wedge \diamond q) \underset{\texttt{HolD}}{\not\Vdash} B \diamond (p \wedge q)$.

4.3 *The challenge precisified*

This suggests a sharpening of the challenge from pernicious transparency (§1): can an assertibility semantics be identified that combines the explanatory strengths of HEx and HolD, at least with respect to the key logical properties surveyed above?

> ***Core Challenge:*** Identify an account of $\Vdash$ that helps to explain (away) the intuition that the assertibility logic of ordinary epistemic vocabulary and attitude reports respects the principles in list A, but not those in B:

A. **KB**, **Factivity**, **Attitude Inheritance**

B. **K-Transparency**, **Modal Omniscience**, **Uncertainty**

This is a challenge for expressivists and descriptivists alike. One way to meet it is with a system that validates everything under A and nothing under B. But the injunction to help *explain* our intuitions is intended to be liberal, allowing explanations that marry a proposed semantics with, say, Gricean pragmatics or error theory. We ourselves exploit this possibility, with a system (§6.2) that answers the Core Challenge despite *not* validating Modal Factivity (and so not Factivity in full generality). Instead, our system offers resources for *explaining away* (as a byproduct of the context-sensitivity of epistemic modals) linguistic data that *prima facie* favors Modal Factivity.

5 Fragmented and Topic-sensitive Expressivism

We now describe a novel system that largely answers the Core Challenge, but with instructive limitations. Call it *FaTE*: **F**ragmented **a**nd **T**opic-sensitive **E**xpressivism.

5.1 *FaTE*

Start by refining the model of epistemic reality from §3.2. A *TF frame* has five components: W, T, @, knowledge function $\mathbf{K}$, and belief function $\mathbf{B}$. W and @ are as before. T is a set of possible *topics*; call a subset of T a *subject matter* (denoted $\mathbf{m}$). We now model a proposition, or information state, as a pair $\langle \mathbf{i}, \mathbf{m} \rangle$: an intension $\mathbf{i}$ plus a subject matter $\mathbf{m}$. The first component gives the verification/truth conditions of a proposition; the second fixes what it is about. A proposition is *veridical at* w iff its intension includes w, and *veridical* iff it is veridical at @.

Per fragmentation, an acceptance state is now modeled as a *set* of propositions, called *fragments*. Thus, $\mathbf{K}$ and $\mathbf{B}$ map a world to a set of propositions: $\mathbf{K}(w)$ is Smith's *total knowledge state at* w and $\mathbf{B}(w)$ is Smith's *total belief state at* w. We stipulate that every proposition in $\mathbf{K}(w)$ is veridical at w and that every knowledge fragment is a *type* of belief fragment: $\mathbf{K}(w) \subseteq \mathbf{B}(w)$, for all w.

A *TF model* $\mathcal{T}$ adds a valuation function $\mathtt{v}$ (mapping each *atom* p to an intension, i.e., its verification/truth conditions) and a *topic assignment* $\mathtt{t}$ (mapping each *sentence* φ to a subject matter, i.e., the set of topics φ is about). Altogether, atomic description p is mapped to a proposition $\langle \mathtt{v}(p), \mathtt{t}(p) \rangle$. Intuitively, this models ordinary discourse. A description like 'Smith is a lawyer' has circumstances under which it is true and under which it is false, and topics that it is about: Smith, Smith's profession, whether Smith is a lawyer. Verification/truth conditions and subject matter shouldn't be confused: '$2+2=4$' and 'every swan is a bird' have the same truth conditions (true under every circumstance) but different subject matter (e.g., numbers versus swans).

Our treatment of topics, like possible worlds, is deliberately abstract and non-committal, resting only on the following structural stipulations for admissible TF models.[27]

SM1. $\mathtt{t}(\neg\varphi) = \mathtt{t}(\diamond\varphi) = \mathtt{t}(\varphi)$ and $\mathtt{t}(\varphi \wedge \psi) = \mathtt{t}(\varphi) \cup \mathtt{t}(\psi)$.

SM2. $\mathtt{t}(B\varphi) \subseteq \mathtt{t}(K\varphi)$ and $\mathtt{t}(\varphi) \subseteq \mathtt{t}(K\varphi)$.

SM3. If $\mathtt{t}(\varphi) \subseteq \mathtt{t}(\psi)$ then: $\mathtt{t}(B\varphi) \subseteq \mathtt{t}(B\psi)$ and $\mathtt{t}(K\varphi) \subseteq \mathtt{t}(K\psi)$.

SM4. If $\langle \mathbf{i}_1, \mathbf{m}_1 \rangle$ and $\langle \mathbf{i}_2, \mathbf{m}_2 \rangle$ are in $\mathbf{K}(w)$, then $\mathbf{m}_1 \cap \mathbf{m}_2 = \varnothing$, i.e., fragments of a total knowledge state have disjoint subject matter.

SM1-SM3 strive to model ordinary discourse.[28] Basic logical operations appear to *merge* the subject matter of their constituents (cf. SM1): the conjunction 'Smith is a lawyer and Jones is a lout' is about Smith, Smith's profession, Jones, Jones' character, etcetera; 'Jones isn't a carpenter' and 'Jones might be a carpenter' pronounce on

[27] This follows a well-established tradition [Berto, 2019, Berto and Hawke, 2021, Hawke et al., 2020, Berto and Hawke, 2022, Hawke et al., forthcoming], echoing discussion of 'awareness' in epistemic logic [Schipper, 2015]. A popular approach to modeling subject matter more concretely identifies topics with (sets of) questions, themselves typically modeled as partitions or divisions of logical space [Lewis, 1988, Yalcin, 2011, Yablo, 2014, Hawke, 2018, Plebani and Spolaore, 2021].

[28] For detailed motivation: Yalcin [2011], [Yablo, 2014, pp.18–19, p.42], Hawke [2018], Berto and Hawke [2022], Hawke et al. [forthcoming].

exactly the same topics as 'Jones is a carpenter': Jones, Jones' profession, whether Jones is a carpenter. As 'Ed is soldier and a gentleman' is about Ed, so too (cf. SM2) 'Smith knows Ed is a soldier and a gentleman' is about Ed (and more: it is about Smith). Given that knowledge is a type of belief, 'Smith knows Ed is a soldier and a gentleman' speaks about what 'Smith believes Ed is a soldier and a gentleman' speaks about (e.g., what Smith believes about Ed), and more, while both speak about what 'Smith believes Ed is a gentleman' speaks about (cf. SM1, SM3).

SM4 is a simplifying assumption, delaying a delicate question: how unified should fragmentationists take total belief or knowledge states to be? Knowledge states resist significant dissonance: for example, it seems impossible for Smith to know p might be true while also knowing p is false. This seems to tell against knowledge states being fragmented without constraint (e.g., separate fragments that have divergent amounts of information on the same subject matter). In contrast, total belief states seem to have leeway for significance dissonance (exhibiting, at best, looser rationality constraints – cf. [Cherniak, 1986, Ch1]). We *approximate* all this by taking knowledge on a subject matter as unified (SM4), without similarly constraining belief.

In what follows, we abuse notation to avoid an excess of symbols: for $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$, we write '$\mathbf{s} \subseteq \mathtt{v}(p)$' to mean $\mathbf{i} \subseteq \mathtt{v}(p)$; '$\mathbf{s} \cap \mathtt{v}(p)$' means $\langle \mathbf{i} \cap \mathtt{v}(p), \mathbf{m} \rangle$; '$\mathtt{t}(\varphi) \subseteq \mathbf{s}$' means $\mathtt{t}(\varphi) \subseteq \mathbf{m}$; '$w \in \mathbf{s}$' means $w \in \mathbf{i}$; $\mathbf{s} = \varnothing$ means $\mathbf{i} = \varnothing$; $\mathbf{s} \neq \varnothing$ means $\mathbf{i} \neq \varnothing$.

Definition 7 (FaTE).

For arbitrary $\boldsymbol{s}$, p, φ *and* ψ, *relative to TF model* $\mathcal{T}$*:*

$\boldsymbol{s} \Vdash p$ *iff* $\mathtt{t}(p) \subseteq \boldsymbol{s}$ *and* $\boldsymbol{s} \subseteq \mathtt{v}(p)$
$\boldsymbol{s} \dashv\Vdash p$ *iff* $\mathtt{t}(p) \subseteq \boldsymbol{s}$ *and* $\boldsymbol{s} \cap \mathtt{v}(p) = \varnothing$
$\boldsymbol{s} \Vdash \neg\varphi$ *iff* $\boldsymbol{s} \dashv\Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash \neg\varphi$ *iff* $\boldsymbol{s} \Vdash \varphi$

$\boldsymbol{s} \Vdash \varphi \wedge \psi$ *iff* $\boldsymbol{s} \Vdash \varphi$ *and* $\boldsymbol{s} \Vdash \psi$
$\boldsymbol{s} \dashv\Vdash \varphi \wedge \psi$ *iff* *there are* $\boldsymbol{u}$ *and* $\boldsymbol{v}$ *s.t.* $\boldsymbol{s} = \boldsymbol{u} \cup \boldsymbol{v}$ *and* $\boldsymbol{u} \dashv\Vdash \varphi$ *and* $\boldsymbol{v} \dashv\Vdash \psi$
$\boldsymbol{s} \Vdash \diamond\varphi$ *iff* $\mathtt{t}(\varphi) \subseteq \boldsymbol{s}$ *and* $\boldsymbol{s} \not\dashv\Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash \diamond\varphi$ *iff* $\boldsymbol{s} \dashv\Vdash \varphi$

This explicitly incorporates topic-sensitivity for atomic descriptions and 'might' claims.[29] For example, p is assertible relative to information $\mathbf{s}$ exactly when the subject matter of p is contained in the subject matter of $\mathbf{s}$ and $\mathbf{s}$ leaves only p-worlds uneliminated. Topic-sensitivity percolates to the boolean expressions. For example, $\neg p$ is assertible exactly when the subject matter of p is contained in the subject matter of $\mathbf{s}$ and every p-world is eliminated by $\mathbf{s}$.

Definition 8 (FaTE continued).

[29] For related approaches to topic-sensitive epistemic possibility claims: Yalcin [2011], Rossi and Özgün [2023].

$\boldsymbol{s} \Vdash K\varphi$ *iff* $\mathtt{t}(K\varphi) \subseteq \boldsymbol{s}$ *and* $\forall w \in \boldsymbol{s}$: $\exists \boldsymbol{k} \in \boldsymbol{K}(w)$: $\boldsymbol{k} \Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash K\varphi$ *iff* $\mathtt{t}(K\varphi) \subseteq \boldsymbol{s}$ *and* $\forall w \in \boldsymbol{s}$: $\forall \boldsymbol{k} \in \boldsymbol{K}(w)$: $\boldsymbol{k} \nVdash \varphi$
$\boldsymbol{s} \Vdash B\varphi$ *iff* $\mathtt{t}(B\varphi) \subseteq \boldsymbol{s}$ *and* $\forall w \in \boldsymbol{s}$: $\exists \boldsymbol{b} \in \boldsymbol{B}(w)$: $\boldsymbol{b} \Vdash \varphi$
$\boldsymbol{s} \dashv\Vdash B\varphi$ *iff* $\mathtt{t}(B\varphi) \subseteq \boldsymbol{s}$ *and* $\forall w \in \boldsymbol{s}$: $\forall \boldsymbol{b} \in \boldsymbol{B}(w)$: $\boldsymbol{b} \nVdash \varphi$

Note the fragmentationist twist, e.g., Bp is assertible relative to **s** only if there is a belief fragment that accepts p at every world uneliminated by **s**. Note that for $\mathbf{s} \Vdash K\varphi$ to hold, it must both be that **s** is about φ *and* that, at every world in **s**, Smith has a knowledge fragment that is about φ.

The definitions of $\vee$, $\models$, $\dashv\models$, etcetera, follow §3.2 and §3.3.

Our system lacks resources to express *justified* belief in particular. Appendix B offers an extension and uses it to model modal Gettier cases.

5.2 *Meeting the challenge (almost)*

In this section, $\Vdash$ and $\models$ are as FaTE defines them. First, a sanity check.

Proposition 16 (Descriptive Factivity). $Kp \underset{\texttt{FaTE}}{\models} p$

Proof. Assume $\mathbf{s} \Vdash Kp$. So, $\mathtt{t}(Kp) \subseteq \mathbf{s}$. So, by SM2, $\mathtt{t}(p) \subseteq \mathbf{s}$. Also, $\forall w \in \mathbf{s}$: $\exists \mathbf{k} \in \mathbf{K}(w)$: $\mathbf{k} \Vdash p$. As $w \in \mathbf{k}$ if $\mathbf{k} \in \mathbf{K}(w)$ (we stipulated that knowledge fragments are veridical), it follows that $\forall w \in \mathbf{s}$: $w \in \mathtt{v}(p)$. So, $\mathbf{s} \subseteq \mathtt{v}(p)$. Altogether: $\mathbf{s} \Vdash p$. □

Topic-sensitivity delivers the next two results.

Proposition 17 (Modal Omniscience fails). $p \underset{\texttt{FaTE}}{\not\models} K \diamond p$

Proof. Counterexample: any model $\mathcal{T}$ and veridical **s** such that (i) $\mathtt{t}(p) \subseteq \mathbf{s}$ and (ii) $\mathbf{s} \subseteq \mathtt{v}(p)$, but (iii) $\mathtt{t}(p) \nsubseteq \mathbf{k}$ for all $\mathbf{k} \in \mathbf{K}(@)$. By (i) and (ii): $\mathbf{s} \Vdash p$. By (iii): $\mathbf{s} \nVdash K \diamond p$, as $@ \in \mathbf{s}$ and $\mathbf{k} \nVdash \diamond p$ for all $\mathbf{k} \in \mathbf{K}(@)$. □

Proposition 18 (K-Transparency fails). $\neg K \neg p \underset{\texttt{FaTE}}{\not\models} K \diamond p$

Proof. Counterexample: any model $\mathcal{T}$ such that (i) $\mathbf{s} = \langle \{@\}, T \rangle$ and (ii) $\mathbf{K}(@) = \{\mathbf{k}_@\} = \{\langle \{@\}, \mathbf{m} \rangle\}$, where $\mathtt{t}(p) \nsubseteq \mathbf{m}$. By (ii): $\mathbf{k}_@ \nVdash \neg p$ (as the topic-containment condition isn't met: $\mathtt{t}(\neg p) \nsubseteq \mathbf{k}_@$, by SM1). So, by (i): $\mathbf{s} \Vdash \neg K \neg p$ (note the topic containment constraint is met: $\mathtt{t}(K \neg p) \subseteq \mathbf{s}$). But by (ii): $\mathbf{k}_@ \nVdash \diamond p$ (as the topic-containment condition isn't met: $\mathtt{t}(\diamond p) \nsubseteq \mathbf{k}(@)$, by SM1). So, by (i): $\mathbf{s} \nVdash K \diamond p$. □

Fragmentation delivers a fourth crucial result.

Proposition 19 (KB). $K\varphi \underset{\texttt{FaTE}}{\models} B\varphi$

Proof. Assume $\mathbf{s} \Vdash K\varphi$. So, $\mathtt{t}(K\varphi) \subseteq \mathbf{s}$ and $\forall w \in \mathbf{s}$: $\exists \mathbf{k} \in \mathbf{K}(w)$: $\mathbf{k} \Vdash \varphi$. So, by SM2, $\mathtt{t}(B\varphi) \subseteq \mathbf{s}$ and, as $\mathbf{K}(w) \subseteq \mathbf{B}(w)$ for all w, we have that $\forall w \in \mathbf{s}$: $\exists \mathbf{k} \in \mathbf{B}(w)$: $\mathbf{k} \Vdash \varphi$. So, $\mathbf{s} \Vdash B\varphi$. □

Finally, FaTE preserves core expressivist advantages.

Proposition 20 (Attitude Inheritance). $B(p \wedge \Diamond q) \underset{\texttt{FaTE}}{\Vdash\!\vDash} B \Diamond (p \wedge q)$. *Similarly,* $K(p \wedge \Diamond q) \underset{\texttt{FaTE}}{\Vdash\!\vDash} K \Diamond (p \wedge q)$.

Proof. Assume $\mathbf{s} \Vdash B(p \wedge \Diamond q)$. So, $\mathtt{t}(B(p \wedge \Diamond q)) \subseteq \mathbf{s}$. By SM1, $\mathtt{t}(p \wedge \Diamond q) = \mathtt{t}(\Diamond(p \wedge q))$. So, by SM3: $\mathtt{t}(B(p \wedge \Diamond q)) = \mathtt{t}(B \Diamond (p \wedge q))$. Thus, $\mathtt{t}(B \Diamond (p \wedge q)) \subseteq \mathbf{s}$. Also, by assumption, $\forall w \in \mathbf{s}$: $\exists \mathbf{b}_w \in \mathbf{B}(w)$: $\mathbf{b}_w \Vdash p \wedge \Diamond q$ (so $\mathbf{b}_w$ contains a $p \wedge q$-world). By SM1, if $\mathbf{b}_w \Vdash p \wedge \Diamond q$ then $\mathtt{t}(\Diamond(p \wedge q)) \subseteq \mathbf{b}_w$. So, $\forall w \in \mathbf{s}$: $\exists \mathbf{b}_w \in \mathbf{B}(w)$: $\mathbf{b}_w \Vdash \Diamond(p \wedge q)$. So, $\mathbf{s} \Vdash B \Diamond (p \wedge q)$. □

6 Factivity and Context

FaTE hasn't yet answered our core challenge. One loose end is that it does *not* validate **Modal Factivity**. Counterexample: any TF model $\mathcal{T}$ such that (i) $\mathbf{s} = \langle\{@\}, T\rangle$, (ii) $@ \notin \mathtt{v}(p)$ and (iii) $\mathbf{K}(@) = \{\mathbf{k}_@\} = \{\langle \mathbf{i}, T\rangle\}$ with $\mathtt{v}(p) \cap \mathbf{i} \neq \varnothing$. By (iii): $\mathbf{k}_@ \Vdash \Diamond p$. So, by (i): $\mathbf{s} \Vdash K \Diamond p$. But by (ii): $\mathbf{s} \nVdash \Diamond p$. Thus, we so far lack an explanation for why bare violations of **Modal Factivity** like (14)-(15) in §2.2 sound incoherent. Of course, a simple modification of FaTE enforces **Factivity**:

$$\mathbf{s} \Vdash K\varphi \text{ iff } \mathbf{s} \Vdash \varphi \text{ and } \mathtt{t}(K\varphi) \subseteq \mathbf{s} \text{ and } \forall w \in \mathbf{s} : \exists \mathbf{k} \in \mathbf{K}(w) : \mathbf{k} \Vdash \varphi$$

But, by itself, this is worryingly *ad hoc*. Besides, we soon survey some reasons to *hesitate* in enforcing Modal Factivity.

A second loose end is that FaTE validates **Uncertainty**. *Proof.* Assume (i) $\mathbf{s} \Vdash K(p \vee q)$ and (ii) $\mathbf{s} \Vdash \neg Kp$. By (i), $\mathtt{t}(K(p \vee q)) \subseteq \mathbf{s}$. By SM1, $\mathtt{t}(\Diamond q) \subseteq \mathtt{t}(p \vee q)$. So, by SM3, $\mathtt{t}(K \Diamond q) \subseteq \mathbf{s}$. Further, by (i), $\forall w \in \mathbf{s}$, $\exists \mathbf{k}_w \in \mathbf{K}(w)$ s.t. $\mathtt{t}(q) \subseteq \mathbf{k}_w$ (by SM1) and every world in $\mathbf{k}_w$ is a $p \vee q$-world. By (ii), $\forall w \in \mathbf{s}$, $\forall \mathbf{k} \in \mathbf{K}(w)$: $\mathbf{k}$ contains a $\neg p$-world. Thus, $\forall w \in \mathbf{s}$, $\exists \mathbf{k}_w \in \mathbf{K}(w)$ s.t. $\mathtt{t}(q) \subseteq \mathbf{k}_w$ and $\mathbf{k}_w$ contains a q-world. So, $\forall w \in \mathbf{s}$: $\exists \mathbf{k}_w \in \mathbf{K}(w)$: $\mathbf{k}_w \Vdash \Diamond q$. Altogether, $\mathbf{s} \Vdash K \Diamond q$.

FaTE, one may check, also delivers: $K \Diamond p \underset{\texttt{FaTE}}{\Vdash\!\vDash} K(p \vee \neg p) \wedge \neg K \neg p$. Combining this with **Uncertainty** validates:

$$\textbf{Restricted K-Transparency: } K(p \vee \neg p) \wedge \neg K \neg p \dashv\Vdash\vDash K \Diamond p$$

But this is another objectionable form of transparency. Consider:

(31) I do not know whether the late Antarctic spring might be caused by ozone depletion. [Yalcin, 2007, pg.1013]

It isn't hard to imagine a context where (31) is perfectly coherent. But, presumably, (31) entails:

(32) I know that the Antarctic spring is or is not caused by ozone depletion.

(33) I don't know that the Antarctic spring is not caused by ozone depletion.

(34) I don't know that the Antarctic spring might be caused by ozone depletion.

In this case, we have an intuitive counterexample to **Restricted K-Transparency**.

So, to meet the Core Challenge of §4.3, a FaTE proponent must satisfactorily explain (away) the intuitive pull of **Modal Factivity** and repulsiveness of **Uncertainty**, either with a well-motivated refinement of FaTE, or pragmatics. I propose that it is advantageous to exploit *both* options, with an independently motivated refinement of FaTE that renders 'might' claims explicitly sensitive to context. First, we consider a tricky question more carefully: is invalidating **Modal Factivity** defensible?

6.1 *Some evidence against modal factivity*

Out of the blue, the following have an air of incoherence:

(35) # It's not raining but Smith knows it might be raining. [Moss, 2018, pg.122]

(36) # Suppose that it isn't raining but Smith knows it might be raining.

Together, **Modal Factivity**, **Inheritance**, and **Attitude Inheritance** explain this effect. Deploying **Modal Factivity**, (35) and (36) entail:

(37) # It's not raining and it might be raining.

(38) # Suppose that it isn't raining and it might be raining.

Deploying **Inheritance** and **Attitude Inheritance**, (37) and (38) entail the obviously incoherent:

(39) # It might both be raining and not raining.

(40) # Suppose that it might both be raining and not raining.

If this is the best explanation, HEx and FaTE (which reject **Modal Factivity**), and HoLD (which rejects **Inheritance** and **Attitude Inheritance**), all face a problem.

However, numerous reasons counsel hesitation in embracing **Modal Factivity** in response. First, related effects in natural language resist an explanation in terms of entailment. Out of the blue, the following sounds terrible in ordinary discourse.

(41) # It's not raining and Smith doesn't know that it's raining.

It is rash to conclude that 'a doesn't know that φ' entails 'φ': when I say 'you don't know that' in response to your worried statement 'the Mets will lose the game', surely I didn't commit myself to the Mets losing. A better explanation of (41) appeals to *pragmatics*, per Stalnaker [1974].

Second, combining **Modal Factivity** with seemingly benign principles yields paradox [Hawke, forthcominga, Sect 7.2]. Consider this abstract chain of reasoning:

(42) $K\neg p \vee Kp \Vdash K\neg \diamond p \vee K\neg \diamond \neg p \Vdash \neg \diamond p \vee \neg \diamond \neg p$

Assuming that constructive dilemma is valid (at least when the disjuncts are knowledge ascriptions), the first implication is delivered by any account that validates an intuitive principle: knowing φ entails knowing that it's not that it might be that $\neg\varphi$ (knowing the train is late entails knowing that it's not that the train might be on time). The second implication is delivered by any account that validates **Modal Factivity**. But applying contraposition and De Morgan's laws to (42) yields:

(43) $\diamond p \wedge \diamond \neg p \Vdash \neg(K\neg p \vee Kp)$

This seems to say, absurdly, that if a speaker in a fixed context is uncertain about p (aptly saying 'it might be and might not be that p'), they can conclude that *Smith* (indeed, *any* agent) is uncertain about p. Rejecting **Factivity** may be counter-intuitive, but this paradox cannot be escaped without rejecting *some* intuitive principle.

Third, sentences like (35) can be rendered palatable by elaborating on the context.

> *Context at time t_1:* You flip a coin and cover the outcome with your hand. Smith saw all this, but (like you) didn't see whether the coin landed Heads or Tails. You both know that the coin flip was fair.

Outlining the situation to Jones (Smith can't hear you), you aptly say:

(44) It might be Heads.

(45) It might be Tails.

(46) Smith knows it might be Heads.

(47) Smith knows it might be Tails.

> *Context at time t_2:* Making sure that Smith cannot see, you take a peek: the coin landed Tails.

Given your new information, you can no longer aptly say (44) to Jones. For *pragmatic* reasons, you shouldn't say (45) either. You should rather say:

(48) It isn't Heads.

(49) It's Tails.

However, it still seems perfectly apt to say both (46) and (47) to Jones. Certainly, it would be odd to say 'Smith knows it might be Tails, but *doesn't* know it might be Heads'. If Smith declines to bet her life-savings on Tails, saying (46) remains a convincing explanation of her praiseworthy behavior. Anyway, it would be puzzling to alter any claim about Smith's knowledge (about the coin) in the transition from t_1 to t_2. Smith neither gains nor loses any information about the coin. Only your information has changed. In the abstract, it would be surprising if what Smith knows *(merely) about the coin* is somehow dependent on what *you* learn about the coin.

As (46) and (48) are both assertible at t_2 (i.e., a fixed context), we apparently have a counterexample to **Modal Factivity**. The following sounds, to my ears, intelligible

in your mouth in context (cf. (35)):

(50) The coin didn't land Heads but all Smith knows is that it *might* be Heads and *might* be Tails.

Fourth, we can identify further attitude attributions that seem factive for straightforward descriptions, but not claims with epistemic vocabulary. Consider the verbs 'forgets' and 'realizes':

(51) # Ed doesn't speak French, but Jen forgot that he does.

(52) # Ed doesn't speak French, but Jen realized that he does.

Now, consider:

> *Context:* Smith forgot to buy cat food. As the store is closed, she panics. She then remembers that you (her neighbor) have many cats and generally have an excess of cat food. She heads over to ask if you can share. Unfortunately, you have run out (and you know it).

Jones asks 'what did Smith want at this odd hour?'. An unobjectionable answer:

(53) She's out of cat food and had at first forgotten that I likely have spare. After realizing that I might have spare food to share, she immediately came to inquire.

But you can only aptly say (54) to Jones, not (55) or (56):

(54) Alas, I don't have any spare cat food.

(55) # I likely have spare cat food.

(56) # I might have spare cat food.

In total, the situation is murky, but it is hardly an *obvious* disadvantage if an expressivist rejects **Modal Factivity**, so long as they answer a residual puzzle: why do assertions like (35) typically sound incoherent? I next argue that expressivists have explanatory tools at hand: an independently motivated contextualist refinement of FaTE predicts that **Modal Factivity** fails in general, but is respected in certain (plausibly prominent) contexts.

6.2 *Contextualist FaTE*

Arguably, linguistic data on epistemic modals motivates a context-sensitive semantics. Descriptivists typically exploit this [Silk, 2015], but expressivists are equally obliged to react. Start with a telling example, quoted from Dorr and Hawthorne.

> [S]uppose you draw a coin from a bucket containing some normal coins and some double-headed coins. Without looking at the coin, you say

(57) I'm not sure whether this coin might land Tails

> Here 'This coin might land Tails' cries out for an interpretation where it is true if the coin is normal and false if it is double-headed. It would be odd to respond to this utterance of [(57)] with [(58)]:

(58) It's obvious that the coin might land Tails, since for all we know it's a normal coin, and normal coins often land Tails when they are tossed

> On the other hand, if after drawing the coin you just utter

(59) This coin might land Tails

> there are two things you could be saying. You could be making the flat-footed remark for which [(58)] would be a reasonable justification, or you could be making a risky guess which would be false if it were a double-headed coin. But either way, in asserting [(59)] you are conveying your lack of knowledge about how the coin will land. It would be misleading to characterize either use as straightforwardly non-epistemic [Dorr and Hawthorne, 2013, pg.13, labeling altered].

The key judgments here can be framed in terms of assertibility, eschewing truth talk. Call the context where (57) is assertible **Coin 1**; call the context where (58) is assertible **Coin 2**. In **Coin 1**, aptly saying (59) is subject to a contextually-determined constraint: the speaker's information must have ruled out the possibility that the coin is double-headed. If her assertions are based on knowledge, the speaker must *know* that the coin is not double-headed. Relatedly, (60) below is assertible in **Coin 1** only if (61) is, and (62) is assertible only if (63) is.

(60) Smith knows the coin might land Tails.

(61) Smith knows the coin is not double-headed.

(62) Smith believes the coin might land Tails.

(63) Smith believes the coin is not double-headed.

Should we conclude that in **Coin 1**, *no* relevant claim of the form 'it might be that p' is assertible? No, for consider:

(64) The coin might land Heads.

The speaker in **Coin 1** can aptly assert (64), despite not having ruled out a double-headed coin. An intuitive explanation: as landing Heads is a possibility in *both* of the salient possible situations – normal coin, double-headed coin – the speaker need not have ruled out the latter to rightly assert (64).

This suggests the following picture. In contexts like **Coin 1**, 'might' claims are regulated by a set of *relevant alternatives*, partitioning the space of possible worlds into distinct, mutually exhaustive cells (in **Coin 1**, the first cell is the set of normal-worlds; the second is the set of double-headed-worlds – ignoring, for simplicity, worlds where the coin doesn't exist, etcetera).[30] Here, the assertibility of might-p requires

[30] For more on modeling relevant alternatives in this manner: [Hawke, 2016a,b]. For more on

not only that the speaker's information is compatible with p, but that it rules out the relevant alternatives incompatible with p (in **Coin 1**, the double-headed situation, when p is 'the coin lands Tails'). In other contexts (e.g. **Coin 2**), no non-trivial partition is deployed: the speaker's information need only be compatible with p.

We refine FaTE to accommodate this. A *TF model in context* (written $\mathcal{T} + \pi$) is a TF model supplemented with a partition π of W, intuitively fixed by context. We call π the *salient distinction(s)* and each of its cells a *relevant alternative*.[31] *Contextualist FaTE* emulates FaTE in its evaluation of formulas (relative to a TF model in context), with the exception:

$\mathbf{s} \Vdash \Diamond\varphi$ iff $\mathtt{t}(\varphi) \subseteq \mathbf{s}$ and $\forall c \in \pi$: if $\mathbf{s} \cap c \neq \varnothing$ then $\mathbf{s} \cap c \not\dashv\Vdash \varphi$
$\mathbf{s} \dashv\Vdash \Diamond\varphi$ iff $\mathtt{t}(\varphi) \subseteq \mathbf{s}$ and $\forall c \in \pi$: if $\mathbf{s} \cap c \neq \varnothing$ then $\mathbf{s} \cap c \Vdash \varphi$

(Again, $\mathbf{s} \cap c$ is shorthand for $\langle \mathbf{i} \cap c, \mathbf{m} \rangle$, where $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$; and $\mathbf{s} \neq \varnothing$ is shorthand for $\mathbf{i} \neq \varnothing$, where $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$.)

In a natural model of **Coin 1**, we have $\pi = \{Normal, Double\}$, $\mathtt{t}(Tails) \subseteq \mathbf{s}$, and $\mathbf{s} \cap Double \neq \varnothing$, where $Normal$ is the set of normal coin worlds, $Double$ is the set of double-headed coin worlds, $Tails$ expresses 'the coin lands Tails', and $\mathbf{s}$ is the speaker's information. Hence, $\mathbf{s} \Vdash \Diamond Tails$ only if $\mathbf{s} \cap Double \not\dashv\Vdash Tails$. But as $Tails$ is false in every world in $Double$, we have $\mathbf{s} \cap Double \dashv\Vdash Tails$.

In a natural model of **Coin 2**, π is the trivial partition $\{W\}$. The clauses for $\Diamond\phi$ become:

$\mathbf{s} \Vdash \Diamond\varphi$ iff $\mathtt{t}(\varphi) \subseteq \mathbf{s}$ and if $\mathbf{s} \neq \varnothing$ then $\mathbf{s} \not\dashv\Vdash \varphi$
$\mathbf{s} \dashv\Vdash \Diamond\varphi$ iff $\mathtt{t}(\varphi) \subseteq \mathbf{s}$ and if $\mathbf{s} \neq \varnothing$ then $\mathbf{s} \Vdash \varphi$

As the speaker's information $\mathbf{s}$ is non-empty in **Coin 2** ($\mathbf{s}$ includes a $Tails$ world), the clauses are here identical to those for FaTE. Thus, $\mathbf{s} \models \Diamond Tails$. More generally, when $\pi = \{W\}$ and $\mathbf{s}$ is veridical (i.e., $@ \in \mathbf{s}$), the clauses for $\Diamond p$ emulate FaTE.

Straightforward modifications of the proofs from §5.2 verify that Contextualist FaTE retains FaTE's key advantages, validating **Descriptive Factivity**, **KB**, and **Attitude Inheritance**, and invalidating **Modal Omniscience** and **K-Transparency**. Further, it continues to invalidate **Modal Factivity**.

However, in any context, a certain distinction is primed for salience: the distinction between *those possible worlds that the speaker's relevant knowledge rules out, and those that it leaves open*. Thus, there is plausibly a prominent (perhaps *default*) class of contexts where this partition is the salient distinction. For example, consider a context (call it **Coin 3**) where we have drawn a coin from a bag and it is highly salient that the speaker knows whether the coin is normal or double-headed. Suppose that in **Coin 3** 'might' (quite naturally) receives a reading akin to **Coin 1** (so 'Smith knows

relevant alternatives for epistemic vocabulary, in general: [Holliday, 2015], [Hawke, forthcomingb].

[31] Theorists that, like Lewis [1988] and Yablo [2014], identify subject matters with divisions of logical space may take π to represent a *topic under discussion*.

the coin might land Tails' implies 'Smith knows the coin is normal' when uttered by the speaker). In such a context, Contextualist FaTE predicts $\neg p \wedge K \diamond p$ is incoherent: the assertibility of 'Smith knows that the coin might land Tails' implies the assertibility of 'The coin might lands Tails'. To see this, take $\Vdash$ defined for Contextualist FaTE, and consider:

Definition 9 (**Diagonal Consequence**). *We write $\varphi \rhd \psi$ to mean: for any $\mathcal{T} + \pi$ and veridical information state **s**, if $\mathbf{s} \Vdash \varphi$ where π is the binary partition $\{\mathbf{i}, \bar{\mathbf{i}}\}$ where $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$ ($\bar{\mathbf{i}}$ being the complement of **i**), then $\mathbf{s} \Vdash \psi$.*

Proposition 21 (**Diagonal Modal Factivity**). $K \diamond p \rhd \diamond p$

Proof. Assume that $\mathbf{s} \Vdash K \diamond p$, where $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$, with $@ \in \mathbf{i}$ and $\pi = \{\mathbf{i}, \bar{\mathbf{i}}\}$. So, $\mathtt{t}(K \diamond p) \subseteq \mathbf{m}$. So, $\mathtt{t}(\diamond p) \subseteq \mathbf{m}$, by SM1 and SM2. Our assumption further implies that $\forall w \in \mathbf{i}$: $\exists \mathbf{k} \in \mathbf{K}(w)$: $\mathbf{k} \Vdash \diamond p$. So, $\forall w \in \mathbf{i}$: $\exists \mathbf{k} \in \mathbf{K}(w)$: $\mathtt{t}(p) \subseteq \mathbf{k}$ and $\forall c \in \{\mathbf{i}, \bar{\mathbf{i}}\}$: if $\mathbf{k} \cap c \neq \varnothing$ then $\mathbf{k} \cap c \not\dashv p$. As $@ \in \mathbf{i}$, we have: $\exists \mathbf{k}_@ \in \mathbf{K}(@)$ such that: if $\mathbf{k}_@ \cap \mathbf{i} \neq \varnothing$ then $\mathbf{k}_@ \cap \mathbf{i} \not\dashv p$. As $@ \in \mathbf{k}_@$, we have $\mathbf{k}_@ \cap \mathbf{i} \neq \varnothing$. Thus, $\mathbf{k}_@ \cap \mathbf{i} \not\dashv p$. So, there must be a p-world in **i**. So, $\mathbf{s} \cap \mathbf{i} \neq \varnothing$ (i.e., $\mathbf{i} \neq \varnothing$) and $\mathbf{s} \cap \mathbf{i} \not\dashv p$ (i.e., $\mathbf{i} \not\dashv p$). So, if $\mathbf{s} \cap \mathbf{i} \neq \varnothing$ then $\mathbf{s} \cap \mathbf{i} \not\dashv p$, and if $\mathbf{s} \cap \bar{\mathbf{i}} \neq \varnothing$ (i.e. $\mathbf{i} \cap \bar{\mathbf{i}} \neq \varnothing$, a necessary falsehood) then $\mathbf{s} \cap \bar{\mathbf{i}} \not\dashv p$. So, $\forall c \in \{\mathbf{i}, \bar{\mathbf{i}}\}$: if $\mathbf{s} \cap c \neq \varnothing$ then $\mathbf{s} \cap c \not\dashv p$. Altogether: $\mathbf{s} \Vdash \diamond p$. □

In short, Contextualist FaTE predicts that **Modal Factivity** holds in a prominent but restricted set of contexts where $\diamond p$ receives a reading akin to **Coin 1** or **Coin 3** and the salient distinction π is $\{\mathbf{i}, \bar{\mathbf{i}}\}$, where $\mathbf{s} = \langle \mathbf{i}, \mathbf{m} \rangle$ is the knowledge the speaker bases her assertions on. Meanwhile, further contexts provide a counterexample to **Modal Factivity** in fully generality (e.g., some contexts where $\diamond p$ receives a reading akin to **Coin 2**).

Finally, Contextualist FaTE, to its credit, invalidates **Uncertainty** and **Restricted K-Transparency**.

Proposition 22 (**Restricted K-Transparency fails**). *For Contextualist FaTE:*

$$K(p \vee \neg p) \wedge \neg K \neg p \not\Vdash K \diamond p$$

Proof. For a counterexample, consider $\mathcal{T} + \pi$ where: (i) $W = \{w_1, w_2, w_3\}$ with $w_1 = @$; (ii) $\mathtt{v}(p) = \{@\} = \{w_1\}$; (iii) $\mathbf{K}(@) = \{\mathbf{k}_@\} = \{\langle \{w_1, w_2\}, \mathtt{t}(p) \rangle\}$; and (iv) $\pi = \{\{w_1, w_3\}, \{w_2\}\}$. Let $\mathbf{s} = \{\{@\}, T\}$. So, $\mathbf{s} \Vdash K(p \vee \neg p) \wedge \neg K \neg p$. However, $\mathbf{s} \not\Vdash K \diamond p$. To see this, note that $\mathbf{k}_@ \cap \{w_2\} = \langle \{w_2\}, \mathtt{t}(p) \rangle$. So, $\mathbf{k}_@ \cap \{w_2\} \neq \varnothing$ (i.e., $\{w_2\} \neq \varnothing$) but $\mathbf{k}_@ \cap \{w_2\} \dashv p$. □

More intuitively: if $K(p \vee \neg p)$ holds at @ (in light of Smith's knowledge fragment $\mathbf{k}_@$ supporting $p \vee \neg p$), but $K \neg p$ and $K \diamond p$ don't hold, the latter's failure cannot be explained as Smith lacking a knowledge fragment that is compatible with p. There must instead be a relevant alternative $c \in \pi$ in play, forcing $\mathbf{k}_@ \not\Vdash \diamond p$: (i) Smith's information fails to eliminate c (i.e., $\mathbf{k}_@ \cap c \neq \varnothing$) yet (ii) Smith would be positioned to deny p if she were to learn c holds (i.e., $\mathbf{k}_@ \cap c \dashv p$).

7 Conclusion

To summarize: first, §4.3 framed a precise version of the challenge from pernicious transparency for assertibility semantics, incorporating logical subtleties and challenging expressivists and descriptivists alike. Second, §5 introduced FaTE, a topic-sensitive and fragmented expressivist semantics that largely addresses the challenge, but with loose ends: **Modal Factivity** isn't validated (without an obvious framework-specific explanation for why it can *appear* valid) and **Uncertainty** (plus **Restricted K-Transparency**) is validated. Third, §6 refined FaTE with independently motivated contextualist machinery, yielding Contextualist FaTE (predicting that our preliminary FaTE system emerges only in certain basic contexts). Contextualist FaTE invalidates **Restricted K-Transparency** and respects **Modal Factivity** in a prominent but restricted class of contexts, offering a tentative explanation for the equivocal ordinary linguistic data on **Modal Factivity**. We conclude: Contextualist FaTE is a promising solution to the precisified challenge from pernicious transparency.

Acknowledgments

Thanks to my reviewers for suggestions that improved the paper. This paper was presented at the Lingnan philosophy department's WIP seminar in Hong Kong SAR (February 20, 2023) and at an APA committee session on Expressivism, at the 2023 Pacific Division APA Meeting at the Westin St. Francis in San Francisco, CA (April 5-8, 2023). Thanks to Nikolaj Pedersen for organizing the latter. Special thanks to Melissa Fusco, Ben Lennertz, and Dan Marshall for detailed commentary. Thanks also to the audiences for helpful comments, especially Ethan Brauer, Rafael De Clercq, and John MacFarlane. This publication is part of the project 'The Semantics of Knowledge and Ignorance' (ECS project number 23603221), funded by the Early Career Scheme of the University Grants Council (UGC) of Hong Kong.

Appendices

A Safety Semantics Minus Factivity

Beddor and Goldstein [2021] exploit the traditional idea that knowledge is a composite: specifically, belief plus truth plus a condition that renders the belief 'safe enough' (cf. [Williamson, 2000]). Their system validates the general factivity of knowledge ascriptions, virtually by stipulation. Compared to our own proposal, a question of motivation is immediately pertinent: if dropping Holism (with well-motivated context-sensitivity) gives room for a sensible theory of knowledge ascription, why deploy controversial tools like safety merely to defuse pernicious transparency? Indeed, theories that take knowledge to be a conjunction of truth, belief, and further conditions are inevitably controversial, given a history of difficulties (cf. Shope [1983], Williamson

[2000]). If truth is independent of the other conditions, generalized Gettierization looms [Zagzebski, 1994]; if not, there is misleading redundancy. Our own account sidesteps such worries.

Beddor and Goldstein's chief rationale for including a 'truth condition' is to assure **Modal Factivity** (in contrast, safely believing descriptive p entails p is true). §6.1 argued that **Modal Factivity** is disputable. We thus consider a variant of Beddor and Goldstein's account that drops the truth condition, judged as a direct competitor to FaTE for meeting the challenge in §4.3.

We work with language $\mathcal{L}^{\blacklozenge}$, including atoms, boolean connectives, 'might' operator $\diamond$, objective possibility operator $\blacklozenge$, and belief operator B. Read $\blacklozenge\varphi$ as 'it could easily have been that φ'. A *safety model* $\mathcal{S}$ is a triple $\langle W,@,\mathbf{b},\mathbf{i}\rangle$. W is the set of all possible worlds, including @. Each world w assigns a truth value $w(p)$ (0 or 1) to each atomic sentence. Functions **b** and **i** map a possible world to an intension: $\mathbf{b}(w)$ (we write $\mathbf{b}^w$) is the agent's *doxastic state* at w (understood as a set of doxastic alternatives), while $\mathbf{i}(w)$ (we write $\mathbf{i}^w$) is the *worldly information* at w: a set of worlds that intuitively are sufficiently 'nearby' w. We stipulate that $\mathbf{i}^w$ is veridical at w, i.e., $w \in \mathbf{i}^w$.

Definition 10 (Safety Semantics). *Given safety model $\mathcal{S}$, $w \in W$ and $\boldsymbol{s} \subseteq W$:*

$$
\begin{array}{lll}
[\![p]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & w(p) = 1 \\
[\![\neg\varphi]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & [\![\varphi]\!]^{w,\boldsymbol{s}} = 0 \\
[\![\varphi \wedge \psi]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & [\![\varphi]\!]^{w,\boldsymbol{s}} = 1 \textit{ and } [\![\psi]\!]^{w,\boldsymbol{s}} = 1 \\
[\![\diamond\varphi]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & \exists v \in \boldsymbol{s}\textit{: } [\![\varphi]\!]^{v,\boldsymbol{s}} = 1 \\
[\![B\varphi]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & [\![\varphi]\!]^{\boldsymbol{b}^w} = 1 \\
[\![\blacklozenge\varphi]\!]^{w,\boldsymbol{s}} = 1 & \textit{iff} & \exists v \in \boldsymbol{i}^w\textit{: } [\![\varphi]\!]^{v,\boldsymbol{i}^v} = 1
\end{array}
$$

where $[\![\varphi]\!]^{\boldsymbol{s}} = 1$ *iff* $\forall w \in \boldsymbol{s}$*:* $[\![\varphi]\!]^{w,\boldsymbol{s}} = 1$.

So $\blacklozenge\varphi$ is true at $\langle w,\mathbf{s}\rangle$ when there is a world v compatible with the worldly information at w (intuitively, v is 'nearby' w) such that φ is true at $\langle v,\mathbf{i}^v\rangle$. To capture assertibility, we take $\mathbf{s} \Vdash \varphi$ to mean $[\![\varphi]\!]^{\mathbf{s}} = 1$. Then, $\varphi \Vdash \psi$ holds when: for every safety model $\mathcal{S}$ and intension $\mathbf{s}$, if $[\![\varphi]\!]^{\mathbf{s}} = 1$ and $@ \in \mathbf{s}$ then $[\![\psi]\!]^{\mathbf{s}} = 1$.

We define $\blacksquare\varphi := \neg\blacklozenge\neg\varphi$ and $K\varphi := B\varphi \wedge \blacksquare(B\varphi \supset \varphi)$. Thus, $\blacksquare(B\varphi \supset \varphi)$ operates as the 'safety condition'. Routine proofs establish that safety semantics then validates **KB** and **Descriptive Factivity**, without validating **K-Transparency**, **Modal Omniscience**, or **Modal Factivity**. However, our safety theory also has some problematic features.

Proposition 23. *Per safety semantics,* ***Attitude Inheritance*** *fails:* $K(p \wedge \diamond q) \nVdash K\diamond(p \wedge q)$.

Proof. Let $\mathcal{S}$ be a safety model: (i) $W = \{w_1,w_2\}$ with $@ = w_1$, (ii) $w_1(p) = w_1(q) = 1$, (iii) $w_2(p) = w_2(q) = 0$, (iv) $\mathbf{b}^{w_1} = \{w_1\}$, (v) $\mathbf{b}^{w_2} = W$, (vi) $\mathbf{i}^{w_1} = W$, and (vii) $\mathbf{i}^{w_2} = \{w_2\}$. Then $[\![B(p \wedge \diamond q)]\!]^{\{w_1\}} = 1$, as there are only $p \wedge q$-worlds in

$\mathbf{b}^{w_1}$. Further, $[\![\blacksquare(B(p \wedge \diamond q) \supset (p \wedge \diamond q))]\!]^{\{w_1\}} = 1$, as for every $v \in \mathbf{i}^{w_1}$, if there is a q-world and only p-worlds in $\mathbf{b}^v$ (note that w_1 meets this condition, but not w_2), then v is a p-world and there is a q-world in $\mathbf{i}^v$. Altogether: $[\![K(p \wedge \diamond q)]\!]^{\{w_1\}} = 1$. However, there exists $v \in \mathbf{i}^{w_1}$ (namely, w_2) such that: there is a $p \wedge q$-world in $\mathbf{b}^v$ (namely, w_1) but no $p \wedge q$-world in $\mathbf{i}^v$ (as only w_2 is in $\mathbf{i}^{w_2}$). Thus, $[\![\blacksquare(B \diamond (p \wedge q) \supset \diamond(p \wedge q))]\!]^{\{w_1\}} = 0$. Thus, $[\![K \diamond (p \wedge q)]\!]^{\{w_1\}} = 0$. □

Thus, our safety theory misses a key advantage of expressivist frameworks like FaTE (cf. §4.1): by itself, it lacks resources to answer an important element of the precisified challenge from transparency (cf. §4.3).

This isn't the end of its problematic logical features.

Proposition 24. *Per safety semantics,* $K \diamond (p \wedge q) \not\Vdash K \diamond p$.

Proof. Let $\mathcal{S}$ be a safety model: (i) $W = \{w_1, w_2, w_3\}$ with $w_1 = @$, (ii) $w_1(p) = w_3(p) = w_1(q) = w_2(q) = 1$, (iii) $w_2(p) = w_3(q) = 0$, (iv) $\mathbf{b}^{w_1} = \{w_1\}$, (v) $\mathbf{b}^{w_2} = \mathbf{b}^{w_3} = \{w_3\}$, (vi) $\mathbf{i}^{w_1} = \{w_1, w_2\}$, and (vii) $\mathbf{i}^{w_2} = \{w_2\}$. As there is a $p \wedge q$-world in $\mathbf{b}^{w_1}$, we have $[\![B \diamond (p \wedge q)]\!]^{\{w_1\}} = 1$. Further, $\forall w \in \mathbf{i}^{w_1}$, if there is a $p \wedge q$-world in $\mathbf{b}^w$ then there's one in $\mathbf{i}^w$. So, $[\![\blacksquare(B \diamond (p \wedge q) \supset \diamond(p \wedge q))]\!]^{\{w_1\}} = 1$. Thus, $[\![K \diamond (p \wedge q)]\!]^{\{w_1\}} = 1$. However, by (v) and (vii), there is a p-world in $\mathbf{b}^{w_2}$ but not in $\mathbf{i}^{w_2}$. So, $\exists w \in \mathbf{i}^{w_1}$ s.t. there's a p-world in $\mathbf{b}^w$ but not in $\mathbf{i}^w$. So, $[\![\blacksquare(B \diamond p \supset \diamond p)]\!]^{\{w_1\}} \neq 1$. So, $[\![K \diamond p]\!]^{\{w_1\}} \neq 1$. □

Thus, unlike FaTE, the current theory erroneously predicts that 'Smith knows that it might be cloudy and damp' does not entail 'Smith knows that it might be cloudy'.

B Modal Gettier Cases

Sarah Moss argues that an adequate theory of knowledge ascription should accommodate *modal Gettier cases*.

(65) **Fake Letters.** Alice enters a psychology study with her friend Bert. As part of the study, each participant is given a detailed survey of romantic questions about their friend. After the study is over, each participant is informed of the probability that they find their friend attractive. Several disgruntled lab assistants have started mailing out fake letters, telling nearly every participant that they probably find their friend attractive. Alice happens to receive a letter from a diligent lab assistant. Her letter correctly reports that she probably does find Bert attractive. Alice reads the letter and comes to have high credence that she finds Bert attractive. [Accordingly, she comes to believe that she might find Bert attractive.] [Moss, 2018, pg.103, additional sentence appended]

Given **Fake Letters**, one reasonably judges that Alice might find Bert attractive and justifiably believes that she might find Bert attractive, but she fails to *know* that she might find Bert attractive, as she could easily have been misled. As Alice in fact finds Bert attractive, she also cannot know that she *doesn't* find him attractive.

So **Fake Letters** serves as an intuitive counterexample to **K-Transparency** [Moss, 2013, 2018, Beddor and Goldstein, 2021].

FaTE has resources for modeling such cases. Consider a TF model $\mathcal{T}$ and information state **s** where: (i) **s** is compatible with p (i.e., **s** is partly about p and is consistent with p), (ii) at every world in **s**, the agent's doxastic state at that world contains a fragment that is compatible with p, and (iii) at every world in **s**, the agent's epistemic state at that world has no fragment compatible with p (in particular, no such fragment is about p, i.e., no such fragment has content whose subject matter includes that of p). It follows that $\mathbf{s} \Vdash \diamond p \wedge B \diamond p \wedge \neg K \diamond p \wedge \neg K \neg p$, as required.

To more pointedly model **Fake Letters**, we can extend FaTE to include justified belief operators. Let a *TF model with justification* be a TF model $\mathcal{J}$ supplemented with a function **J** that maps a world to fragments of justified belief: $\mathbf{J}(w)$ is Smith's *total justified belief state at* w. We assume that every justified belief fragment is a *type* of belief fragment ($\mathbf{J}(w) \subseteq \mathbf{B}(w)$, for all w) and every knowledge fragment is a *type* of justified belief fragment ($\mathbf{K}(w) \subseteq \mathbf{J}(w)$, for all w). Again, our semantic treatment remains silent on the epistemological question as to what *makes* a fragment justified. For sensible constraints on subject matter, we assume:

SM5. $\mathtt{t}(B\varphi) \subseteq \mathtt{t}(J\varphi)$ and $\mathtt{t}(J\varphi) \subseteq \mathtt{t}(K\varphi)$.

SM6. If $\mathtt{t}(\varphi) \subseteq \mathtt{t}(\psi)$ then: $\mathtt{t}(J\varphi) \subseteq \mathtt{t}(J\psi)$.

FaTE with Justification: We extend the semantics for FaTE with:

$\mathbf{s} \Vdash J\varphi$ iff $\mathtt{t}(J\varphi) \subseteq \mathbf{s}$ and $\forall w \in \mathbf{s}$: $\exists \mathbf{j} \in \mathbf{J}(w)$: $\mathbf{j} \Vdash \varphi$
$\mathbf{s} \dashv\Vdash J\varphi$ iff $\mathtt{t}(J\varphi) \subseteq \mathbf{s}$ and $\forall w \in \mathbf{s}$: $\forall \mathbf{j} \in \mathbf{J}(w)$: $\mathbf{j} \nVdash \varphi$

Fake Letters can then be modeled with a model $\mathcal{J}$ and proposition **s** with the following features: (i) **s** is compatible with p (i.e., **s** is partly about p and is consistent with p), (ii) at every world in **s**, the agent's justified belief state at that world contains a fragment that is compatible with p, and (iii) at every world in **s**, the agent's knowledge state contains no fragment compatible with p (in particular, no such fragment is about p, i.e., no such fragment has content whose subject matter includes that of p). It follows that $\mathbf{s} \Vdash \diamond p \wedge B \diamond p \wedge J \diamond p \wedge \neg K \diamond p \wedge \neg K \neg p$.

In short, FaTE can diagnose a modal Gettier case with respect to $\diamond p$ as a situation where an agent's cognitive system contains belief fragments about p's subject matter, but no *knowledge* fragments about p's subject matter. This doesn't imply that the agent is unable to grasp p's subject matter, enter into reasoning with content about that subject matter, or attend to the question as to whether p is true: intuitively, these functions could manifest via the agent's *belief* fragments on p's subject matter.

However, a deeper worry points again to FaTE's limitations. If no fragment of her knowledge is about p's subject matter, our agent has no knowledge *at all* on that subject matter. But, intuitively, modal Gettier cases exist where the agent in question has *some* knowledge about the subject matter of p. In **Fake Letters**, it would be odd to deny that Alice at least knows that either she finds Bert attractive or she doesn't ($p \vee \neg p$). By the lights of FaTE, Alice must have a knowledge fragment about p's

subject matter, grounding her knowledge that $p \vee \neg p$. More pointedly, recall (§6) that FaTE validates **Restricted K-Transparency**: $K(p \vee \neg p) \wedge \neg K \neg p \Rightarrow\!\!\!\models K \diamond p$. According to FaTE, if $K(p \vee \neg p)$ and $\neg K \neg p$ hold, our agent *cannot* be in a modal Gettier case, contrary to our intuitions about **Fake Letters**.

There is an answer: shift to Contextualist FaTE, as one of its chief virtues is that it invalidates **Restricted K-Transparency** (§6.2). Hence, Contextualist FaTE offers improved tools for modeling modal Gettier cases, with nuanced *explanatory* options. By its lights, if 'Alice knows $p \vee \neg p$' is true, but 'Alice knows $\neg p$' and 'Alice knows $\diamond p$' are false, there must be a contextually salient distinction in play: there must be a relevant alternative c such that (i) Alice's information fails to eliminate c and (ii) Alice would be positioned to deny p if she were to learn c holds. In **Fake Letters**, there is an obvious candidate for c: the possibility that Alice's survey indicates that she doesn't find Bert attractive. Just as 'Smith knows the coin might land Tails' is true in **Coin 1** (§6.2) only if Smith knows that the coin isn't double-headed, so 'Alice knows she might find Bert attractive' is true in **Fake Letters** only if Alice knows that her survey doesn't indicate that she doesn't find Bert attractive.

References

Maria Aloni. Logic and Conversation: The Case of Free Choice. *Semantics and Pragmatics*, 15(5):565–589, 2022.

Maria Aloni, Luca Incurvati, and Julian J. Schlöder. Epistemic modals in hypothetical reasoning. *Erkenntnis*, 88(8):3551–3581, 2023. doi: 10.1007/s10670-022-00517-x.

Dorit Bar-On, Matthew Chrisman, and James Sias. (how) is ethical neo-expressivism a hybrid view? In Guy Fletcher and Michael Ridge, editors, *Having It Both Ways: Hybrid Theories and Modern Metaethics*, pages 223–247. Oxford University Press, 2014.

Bob Beddor and Simon Goldstein. Mighty Knowledge. *The Journal of Philosophy*, 118(5):229–269, 2021.

Francesco Berto. Aboutness in imagination. *Philosophical Studies*, 175(8):1871–1886, 2018.

Francesco Berto. Simple hyperintensional belief revision. *Erkenntnis*, 84(3):559–575, 2019.

Francesco Berto and Peter Hawke. Knowability relative to information. *Mind*, 130 (517):1–33, 2021. doi: 10.1093/mind/fzy045.

Francesco Berto and Peter Hawke. Two-component semantics. In *Topics of Thought: The Logic of Knowledge, Belief, Imagination*, pages 21–59. Oxford University Press, 2022.

Francesco Berto and Aybüke Özgün. The logic of framing effects. *The Journal of Philosophical Logic*, 52:939–962, 2023.

Justin Bledin. Logic Informed. *Mind*, 123(490):277–316, 2014. doi: 10.1093/mind/fzu073.

Justin Bledin and Tamar Lando. Closure and Epistemic Modals. *Philosophy and Phenomenological Research*, 97(1):3–22, 2018.

Cristina Borgoni, Dirk Kindermann, and Andrea Onofri, editors. *The Fragmented Mind*. Oxford University Press, 2021.

John Cantwell. Unity and autonomy in expressivist logic. *Dialectica*, 68(3):443–457, 2014. doi: 10.1111/1746-8361.12072.

John Cantwell. An expressivist analysis of the indicative conditional with a restrictor semantics. *Review of Symbolic Logic*, 14(2):487–530, 2021. doi: 10.1017/s1755020319000662.

Christopher Cherniak. *Minimal Rationality*. Bradford Books: MTT Press, 1986.

Matthew Chrisman. From epistemic contextualism to epistemic expressivism. *Philosophical Studies*, 135(2):225–254, 2007. doi: 10.1007/s11098-005-2012-3.

Ivano Ciardelli. Restriction without Quantification: Embedding and Probability for Indicative Conditionals. *Ergo*, 8, 2021.

Cian Dorr and John Hawthorne. Embedding Epistemic Modals. *Mind*, 122(488): 867–913, 2013. doi: 10.1093/mind/fzt091.

Janice L Dowell. A Flexible Contextualist Account of Epistemic Modals. *Philosophers' Imprint*, 11(14):1–25, 2011.

Paul Egré. Question-embedding and factivity. *Grazer Philosophische Studien*, 77(1): 85–125, 2008.

Ronald Fagin and Jospeh Halpern. Belief, awareness and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

Kit Fine. A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic*, 46(6):625–674, 2017.

Christopher Gauker. *Conditionals in Context*. MIT Press, 2005.

Allan Gibbard. *Thinking How to Live*. Harvard University Press, 2003.

Daniel Greco. Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3):656–673, 2015.

Peter Hawke. Relevant alternatives in epistemology and logic. In Ángel Nepomuceno Fernández, Olga Pombo Martins, and Juan Redmond, editors, *Epistemology, Knowledge and the Impact of Interaction*. Springer Verlag, 2016a.

Peter Hawke. Questions, topics and restricted closure. *Philosophical Studies*, 173 (10):2759–2784, 2016b. doi: 10.1007/s11098-016-0632-4.

Peter Hawke. Theories of aboutness. *Australasian Journal of Philosophy*, 96(4): 697–723, 2018. doi: 10.1080/00048402.2017.1388826.

Peter Hawke. Stable acceptance for mighty knowledge. *Philosophical Studies*, forthcominga.

Peter Hawke. Relevant alternatives and missed clues: redux. *Journal of Philosophy*, forthcomingb.

Peter Hawke and Shane Steinert-Threlkeld. Informational Dynamics of 'Might' Assertions. In Wiebe van der Hoek, Wesley H Holliday, and Wen-fang Wang, editors, *Proceedings of Logic, Rationality, and Interaction (LORI-V)*, volume 9394 of *Lecture Notes in Computer Science*, pages 143–155, 2015.

Peter Hawke and Shane Steinert-Threlkeld. Informational dynamics of epistemic possibility modals. *Synthese*, 195(10):4309–4342, 2018. doi: 10.1007/s11229-016-1216-8.

Peter Hawke and Shane Steinert-Threlkeld. Semantic expressivism for epistemic modals. *Linguistics and Philosophy*, 44(2):475–511, 2021. doi: 10.1007/s10988-020-09295-7.

Peter Hawke, Aybüke Özgün, and Francesco Berto. The fundamental problem of logical omniscience. *Journal of Philosophical Logic*, 49(4):727–766, 2020. doi: 10.1007/s10992-019-09536-6.

Peter Hawke, Levin Hornischer, and Francesco Berto. Truth, topicality, and transparency: One-component versus two-component semantics. *Linguistics and Philosophy*, forthcoming.

Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, 1962.

D. Hoek. Questions in action. *Journal of Philosophy*, 119(3):113–143, 2022.

Daniel Hoek. Minimal Rationality and the Web of Questions. In Andy Egan, Peter van Elswyk, and Dirk Kindermann, editors, *Unstructured Content*. Oxford University Press, forthcoming.

Wesley H. Holliday. Epistemic closure and epistemic logic i: Relevant alternatives and subjunctivism. *Journal of Philosophical Logic*, 44(1):1–62, 2015. doi: 10.1007/s10992-014-9338-2.

Wesley H. Holliday and Matthew Mandelkern. The orthologic of epistemic modals. *The Journal of Philosophical Logic*, forthcoming.

Luca Incurvati and Julian Schlöder. Weak assertion. *The Philosophical Quarterly*, 69(277):741–777, 2019.

Luca Incurvati and Julian Schlöder. *Reasoning With Attitude*. Oxford University Press, 2023.

Frank Jackson. *Conditionals*. Oxford University Press, 1987.

Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39:341–50, 1984.

Angelika Kratzer. The Notional Category of Modality. In Hans-Jürgen Eikmeyer and Hannes Rieser, editors, *Words, Worlds, and Context*, pages 38–74. Walter de Gruyter, 1981.

Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, 2012.

H. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2000.

David Lewis. Logic for Equivcators. *Noûs*, 16(3):431–441, 1982.

David Lewis. Relevant Implication. *Theoria*, 54(3):161–174, 1988.

Hanti Lin. The Meaning of Epistemic Modality and the Absence of Truth. In Syraya Chin-Mu Yang, Duen-Min Deng, and Hanti Lin, editors, *Structural Analysis of Non-Classical Logics*. Springer, 2016.

John MacFarlane. *Assessment Sensitivity*. Oxford University Press, 2014. doi: 10.1093/acprof:oso/9780199682751.001.0001.

Matthew Mandelkern. Bounded Modality. *The Philosophical Review*, 128(1):1–61, 2019.

Sarah Moss. Epistemology Formalized. *The Philosophical Review*, 122(1):1–43, 2013. doi: 10.1215/00318108-1728705.

Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015. doi: 10.3765/sp.8.5.

Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, 2018.

Matteo Plebani and Giuseppe Spolaore. Subject matter: A modest proposal. *Philosophical Quarterly*, 71(3):605–622, 2021. doi: 10.1093/pq/pqaa054.

Vít Punčochář. Weak Negation in Inquisitive Semantics. *Journal of Logic, Language and Information*, 24(3):323–355, 2015. doi: 10.1007/s10849-015-9219-2.

Vít Punčochář. A Generalization of Inquisitive Semantics. *Journal of Philosophical Logic*, 45(4):399–428, 2016. doi: 10.1007/s10992-015-9379-1.

Agustín Rayo. *The Construction of Logical Space*. Oxford University Press, 2013.

Michael Ridge. *Impassioned Belief*. New York, NY: Oxford University Press, 2014.

Niccolò Rossi and Aybüke Özgün. A hyperintensional approach to positive epistemic possibility. *Synthese*, 202(44):1–29, 2023. doi: 10.1007/s11229-023-04219-x.

Burkhard C. Schipper. Awareness. In H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi, editors, *Handbook of Epistemic Logic*, pages 79–146. College Publications, 2015.

Mark Schroeder. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford University Press, New York, 2008a.

Mark Schroeder. Expression for Expressivists. *Philosophy and Phenomenological Research*, 76(1):86–116, 2008b. ISSN 00318205. doi: doi:10.1111/j.1933-1592.2007.00116.x.

Mark Schroeder. Hybrid expressivism: Virtues and vices. *Ethics*, 119(2):257–309, 2009.

Mark Schroeder. Two roles for propositions: Cause for divorce? *Noûs*, 47(3): 409–430, 2011. doi: 10.1111/j.1468-0068.2011.00833.x.

Mark Schroeder. Attitudes and Epistemics. In *Expressing Our Attitudes*, pages 225–256. Oxford University Press, 2015.

Robert K. Shope. *The Analysis of Knowledge: A Decade of Research*. Princeton University Press, 1983.

Alex Silk. How to Be an Ethical Expressivist. *Philosophy and Phenomenological Research*, 91(1):47–81, 2015. doi: 10.1111/phpr.12138.

Alex Silk. How to embed an epistemic modal: attitude problems and other defects of character. *Philosophical Studies*, 174(7):1773–1799, 2017. doi: 10.1007/s11098-016-0827-8.

Robert Stalnaker. Pragmatic presuppositions. In M. K. Munitz and Peter Unger, editors, *Semantics and Philosophy*. New York University Press, 1974.

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

Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.

Frank Veltman. *Logics for Conditionals*. PhD thesis, Universiteit van Amsterdam, 1985.

Frank Veltman. Defaults in Update Semantics. *Journal of Philosophical Logic*, 25 (3):221–261, 1996. doi: 10.1007/BF00248150.

Malte Willer. Dynamics of Epistemic Modality. *Philosophical Review*, 122(1):45–92, 2013. doi: 10.1215/00318108-1728714.

Malte Willer. Simplifying with Free Choice. *Topoi*, 2017a. doi: 10.1007/s11245-016-9437-5.

Malte Willer. Advice for Noncognitivists. *Pacific Philosophical Quarterly*, 98(2017): 174–207, 2017b. doi: 10.1111/papq.12160.

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

Stephen Yablo. *Aboutness*. Princeton University Press, 2014.

Seth Yalcin. Epistemic Modals. *Mind*, 116(464):983–1026, 2007. doi: 10.1093/mind/fzm983.

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

Seth Yalcin. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, 112(2): 123–160, 2012a. doi: 10.1111/j.1467-9264.2012.00329.x.

Seth Yalcin. Context Probabilism. In Maria Aloni, Vadim Kimmelman, Floris Roelofsen, Galit W Sassoon, Katrin Schulz, and Matthijs Westera, editors, *Logic, Language, and Meaning: Amsterdam Colloquium 2011*, volume 7218 of *Lecture Notes in Computer Science*, pages 12–21. Springer-Verlag, 2012b.

Seth Yalcin. Belief as Question-Sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47, 2018. doi: 10.1111/phpr.12330.

Fan Yang and Jouko Väänänen. Propositional Team Logics. *Annals of Pure and Applied Logic*, 168(7):1406–1441, 2017. doi: 10.1016/j.apal.2017.01.007.

Linda Zagzebski. The inescapability of gettier problems. *The Philosophical Quarterly*, 174(44):65–73, 1994.