Indicative Conditionals and Epistemic Luminosity

Matthew Hewson
Merton College, University of Oxford, UK
matthew.hewson@philosophy.ox.ac.uk

James Ravi Kirkpatrick
Somerville College, University of Oxford, UK
james.kirkpatrick@philosophy.ox.ac.uk

Kevin Dorst has recently pointed out an apparently puzzling consequence of denying epistemic luminosity: given some natural-sounding bridging principles between knowledge, credence, and indicative conditionals, the denial of epistemic luminosity licenses the knowledge and assertability of abominable-sounding conditionals of the form ⌜If I don't know that ϕ, then ϕ⌝. We provide a general and systematic examination of this datum by testing Dorst's claim against various semantics for the indicative conditional in the setting of epistemic logic. Our conclusion is that, regardless of whether knowledge is luminous, the knowability of these conditionals is highly sensitive to the correct semantic analysis of the indicative conditional. Moreover, standard pragmatic resources can explain away the infelicity of such assertions. As it stands, the datum does not tell against epistemic non-luminosity.

1. Introduction

There is a much-discussed view in epistemology according to which one may sometimes know a proposition *p* while failing to be in a position to know that one knows *p*. Then we say that one's knowledge that *p* is *not epistemically luminous*.[1] More generally, this view denies the validity of introspection principles like the following:

Strong KK. For any proposition *p*, if one knows that *p*, then one knows that one knows that *p*.

Weak KK. For any proposition *p*, if one knows that *p*, then one is in a position to know that one knows that *p*.

The case against these introspection principles is compelling. For if knowledge requires that one possesses the relevant concept, then one

[1] We borrow this terminology from Williamson (2000, ch. 4). A state *s* is luminous if and only if whenever one is in *s*, one is in a position to know that one is in *s*.

Mind, Vol. 131 . 521 . January 2022
https://doi.org/10.1093/mind/fzab064 Advance Access publication 13 December 2021

may fail to know whether one knows by failing to possess the concept of knowledge (Feldman 1981). Furthermore, epistemic luminosity is widely held to be incompatible with the possibility of inexact knowledge arising from imprecise measurement and estimation (Williamson 1992, 2000).

However, Kevin Dorst (2019) has recently challenged the sovereignty of epistemic non-luminosity. With the help of some bridging principles between epistemic and credal states and conditionals, Dorst argues that failures of epistemic luminosity entail the knowability and assertability of 'abominable' conditionals. For example, he argues that the non-luminosity of one's knowledge that Padua is in Italy entails that one knows (1):[2]

(1) #If I don't know that Padua is in Italy, then Padua is in Italy.

Not only does (1) sound odd and unassertable at a context, knowing the content of (1) also seems to be an epistemically dubious position for an agent to be in. But Dorst argues that when one's knowledge that *p* is not luminous, one is in a position to know the conditional 'If I don't know that *p*, then *p*', whereas epistemic luminosity provides an explanation of the unknowability and unassertability of such conditionals.

Before accepting Dorst's conclusions, it is worth investigating whether the data can be explained from within a broadly knowledge-centric epistemology, one that does not rely on potentially objectionable hypotheses about the relation between credence and knowledge. More specifically, in this paper, we will use the setting of epistemic logic to check which semantic analyses of the indicative conditional, together with epistemic (non-)luminosity, entail knowledge of abominable conditionals. Furthermore, we shall develop an explanation for why such conditionals are unassertable. We argue that neither defenders nor opponents of epistemic luminosity escape Dorst's charge unscathed; whether one avoids knowledge of abominable conditionals is highly sensitive to which theory of the conditional is correct.

This paper is structured as follows. §2 criticizes Dorst's argument and argues that this issue should be pursued in the setting of epistemic

[2] For other arguments in favour of the KK principle, see Das and Salow (2018), Goodman and Salow (2018), and Greco (2014, 2015).

logic. §3 introduces some structural constraints for when failures of epistemic luminosity entail knowledge of abominable conditionals. §§4–5 discuss how the semantics for the indicative conditional affect whether knowledge of abominable conditionals follows from epistemic non-luminosity and epistemic luminosity, respectively. §6 explains the infelicity of abominable conditionals and related issues concerning belief. §7 concludes.

2. A puzzle about iffy knowledge

Dorst's argument relies on two bridging principles between conditionals, knowledge, and credence:[3]

DILIGENCE. $\Diamond(Kp \wedge \neg KKp) \supset \Diamond(Kp \wedge \neg KKp \wedge \forall q(Kq \leftrightarrow Cr(q) = 1))$

STABILITY. $\Box((Cr(p) > 0 \wedge Cr(q) = 1) \supset Cr(p \rightarrow q) = 1)$

According to DILIGENCE, if epistemic luminosity can fail at all, then it can fail for agents who are *diligent*; that is, agents who are certain of all and only that which they are in a *position to know*.[4] This principle, Dorst tells us, follows from the natural hypothesis that full credence is the 'internal component' of knowledge (Dorst 2019, p. 1228). Moreover, it follows from the natural hypotheses that one should proportion one's credence in line with one's evidence, and that one knows all and only those things that are certain on one's evidence.

According to STABILITY, if you leave open that p while being certain that q, then you are certain that if p, q. Dorst tells us that STABILITY falls out of popular accounts of the indicative conditional.[5]

[3] $\Diamond$ and $\Box$ denote metaphysical possibility and necessity, respectively. *Cr* is a credence function mapping propositions to $[0,1]$; we stipulate that credence functions are probabilistic and defined over a *finite* algebra.

[4] We shall move freely between talk of DILIGENCE (the *principle*) and agents who exemplify DILIGENCE. Strictly speaking, the principle just says that such agents are possible. Frequently, for the purposes of our investigation, we will talk as though they are actual. It should always be clear from the context which we mean.

[5] We revisit this claim in §4.

Given these assumptions and a particular failure of epistemic luminosity, it follows that one can know the corresponding abominable conditional:

(1) $Kp \wedge \neg KKp$ (assumption)

(2) $\Diamond Cr(p) = 1$ (by 1 and Diligence)

(3) $\Diamond Cr(\neg Kp) > 0$ (by 1, $Cr(Kp) < 1$, $\neg K \neg Kp$)

(4) $\Diamond Cr(\neg Kp \rightarrow p) = 1$ (by 2, 3, and Stability)

(5) $\Diamond K(\neg Kp \rightarrow p)$ (by 4 and Diligence)

If this argument is correct, proponents of epistemic non-luminosity should be worried.

The argument is striking and its conclusion is compelling, especially given the intuitive appeal of the bridging principles. However, one may wonder to what extent Dorst's conclusion relies on them and whether similar results may be generated in a more general setting. In the next section, we show that Dorst's bridging principles are inessential to the core of his argument; we provide the same sort of investigation by appealing to different models in epistemic logic and testing whether Dorst's conclusions hold on a variety of semantics for the indicative conditional. This framework does without the credally infallible agents invoked by Diligence, and better fits the broadly knowledge-first approach to theorizing we endorse. Investigating Dorst's claims in this cleaner setting will keep us honest; when our predictions disagree with Dorst's, we will be able to spell out exactly why our respective frameworks differ.

3. Abominable conditionals in epistemic logic

Failures of epistemic luminosity can be modelled using a familiar framework taken from the standard possible worlds semantics for epistemic logic (Hintikka 1962). This framework lets us precisely formulate necessary and sufficient conditions for knowing abominable conditionals when epistemic luminosity fails, and check the consistency and consequences of our cases. For simplicity, we focus only on single-agent cases.[6] With this restriction in place, this section centres on structures known as *models*.

[6] See Williamson (2000, pp. 131–4) for interpersonal cases involving the failure of KK.

Let $\mathcal{L}$ be the formal language under consideration, with the following Backus–Naur form:

$$\mathcal{L} ::= P \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid K\phi \mid \phi \rightarrow \psi$$

where $P \in \textsc{atom}(\mathcal{L})$ are the atomic sentences of the language, $\neg$, $\wedge$, $\vee$ are the usual sentential connectives, $\ulcorner K\phi \urcorner$ formalizes $\ulcorner \alpha$ knows that $\phi \urcorner$, and $\ulcorner \phi \rightarrow \psi \urcorner$ formalizes the natural language indicative conditional $\ulcorner$If ϕ, then $\psi \urcorner$.[7]

A model $\mathcal{M}$ $(= \langle W, R, I \rangle)$ for $\mathcal{L}$ is an ordered triple, consisting of a finite non-empty set W, a binary relation R on W (a set of ordered pairs of members of W), and an interpretation function I mapping sentence letters to truth-values relative to worlds.

Propositional notions are encoded in the model by W. Informally, we may think of W as a set of relevant possible worlds. We may think of subsets of W as propositions. Truth and entailment are understood set-theoretically: a proposition $p \subseteq W$ is true at a world w if and only if $w \in p$; the propositions $p_1, \ldots, p_n$ entail q if and only if $p_1 \cap \ldots \cap p_n \subseteq q$. The usual logical constants are defined by their set-theoretic counterparts: the negation of a proposition is its complement, the conjunction of two propositions is their set-theoretic intersection, the disjunction of two propositions is their set-theoretic union, and so on.

Epistemic notions are encoded in the model by R. Informally, we may think of R as an epistemic accessibility relation on W, which tells us which worlds are 'seen' by which worlds. $R(w)$ is the set of epistemically accessible worlds at w. More specifically, world v is epistemically possible at a world w (Rwv) if and only if, for all one knows in w, one is in v; in other words, whatever one knows in w is true in v. On this basis, we may define a function $K : \wp(W) \mapsto \wp(W)$ as follows, where $p \subseteq W$:[8]

$$Kp = \{w \in W : \forall w' \in W(Rww' \supset w' \in p)\}$$

Informally, Kp is the proposition that one knows that p. That is, one knows that p if and only if p is true at every world consistent with what one knows.[9]

[7] Here we use Quine-corners in the familiar sense.

[8] Here we abuse the distinction between the object language and metalanguage.

[9] While the background possible-worlds framework generates problems of logical omniscience, this idealization is benign in this setting; see Stalnaker (1999) and Williamson (2014).

Given that knowledge is factive, we impose one constraint on the accessibility relation: reflexivity ($\forall wRww$). That is, every world is epistemically possible to itself. Consequently, $Kp \subseteq p$, for any $p \subseteq W$, and $w \in R(w)$, for any w.

Semantic notions are encoded in the model by I. Informally, we may think of I as an interpretation function that assigns values to the atomic sentences in our language. In particular, let $I : \langle \alpha, w \rangle \mapsto \{1, 0\}$ map each atomic sentence of our language to a truth-value relative to a world.

It is then the job of the valuation function of the model to assign values to complex sentences in our language, again relative to worlds. Let $[\![.]\!]^w$ be a valuation function from sentences in $\mathcal{L}$ to truth-values relative to w, the world of evaluation, that satisfies the following constraints:[10]

- $[\![\alpha]\!]^w = I(\alpha, w)$, where α is an arbitrary element of $\textsc{atom}(\mathcal{L})$
- $[\![\ulcorner \neg\phi \urcorner]\!]^w = 1$ if and only if $[\![\phi]\!]^w = 0$
- $[\![\ulcorner \phi \wedge \psi \urcorner]\!]^w = 1$ if and only if $[\![\phi]\!]^w = [\![\psi]\!]^w = 1$
- $[\![\ulcorner \phi \vee \psi \urcorner]\!]^w = 1$ if and only if $[\![\phi]\!]^w = 1$ or $[\![\psi]\!]^w = 1$
- $[\![\ulcorner K\phi \urcorner]\!]^w = 1$ if and only if for all $w' \in W$, if Rww', then $[\![\phi]\!]^{w'} = 1$

Since our central contention is that the relationship between epistemic luminosity and knowledge of abominable conditionals crucially depends on the precise semantics for indicative conditionals, we shall remain non-committal on the semantics of '$\rightarrow$' at this stage. We shall nevertheless observe that any adequate theory of indicative conditionals must predict that a conditional is false at a world whenever its antecedent is true and its consequent is false at that world.

Counterexamples to epistemic luminosity have a well-known structure in these models. Let w be an arbitrary world in which one knows that ϕ but does not know that one knows that ϕ. Given that one knows that ϕ at w, it follows that ϕ is true at every world that is epistemically accessible relative to w. Furthermore, given that one does not know that one knows that ϕ, it follows that there is some epistemically accessible world w' relative to w in which one does not

[10] Strictly speaking, the valuation function would be indexed to a model, but for the purposes of simplicity these indexes are omitted.

know that ϕ; that is, there is some epistemically accessible world w'' relative to w' at which ϕ is false. It is easy to verify that the epistemic accessibility relation of this model must be non-transitive. Indeed, this is a necessary and sufficient condition for luminosity failures (Williamson 2014, p. 975).

For illustration, consider the following toy example $\mathcal{M} = \langle W, R, I \rangle$, where $W = \{w_i \mid \forall i \in \{1, 2, 3\}\}$ and $R = \{\langle w_i, w_j \rangle \mid i = j \text{ or } j = i + 1, \forall i, j \in \{1, 2, 3\}\}$. Note that R is a reflexive relation, but not transitive. R is not transitive, because $\langle w_1, w_2 \rangle \in R$ and $\langle w_2, w_3 \rangle \in R$, but $\langle w_1, w_3 \rangle \notin R$. It is easy to verify that epistemic luminosity fails on this model. Let $[\![\phi]\!] = \{w_1, w_2\}$. Then $[\![\ulcorner K\phi \urcorner]\!] = \{w_1\}$: one knows that ϕ in w_1 because ϕ is true in every epistemically possible world in w_1; one does not know that ϕ in w_2 because ϕ is false at w_3, which is epistemically possible at w_2. Consequently, $[\![\ulcorner KK\phi \urcorner]\!] = \varnothing$: thus one does not know that one knows that ϕ at w_1 because one does not know that ϕ at w_2, which is epistemically possible at w_1. Therefore, epistemic luminosity fails in w_1, since $\ulcorner K\phi \urcorner$ is true but $\ulcorner KK\phi \urcorner$ is false there.[11] Furthermore, without specifying a precise semantics for the indicative conditional, observe that *if* one knows the corresponding abominable conditional at w_1, then the abominable conditional is true at both w_1 and w_2, given the semantics for K. Moreover, the abominable conditional is not true at w_3, since the antecedent is true and the consequent is false at that world. Interestingly, one's knowledge of the abominable conditional at w_1 is not luminous: one fails to know that one knows the conditional. This is because one fails to know the abominable conditional at w_2. Figure 1 illustrates these observations, indicating which sentences are true at which worlds.

This section has outlined some necessary and sufficient conditions for when luminosity failures entail knowledge of abominable conditionals. In particular, we have illustrated the structural restrictions on the semantics for indicative conditionals needed for such a situation to arise. The following section will examine whether any prominent semantic analyses are compatible with these structural requirements.

[11] For compelling real-world failures of epistemic luminosity, see Williamson (2000, ch. 4; 2014).

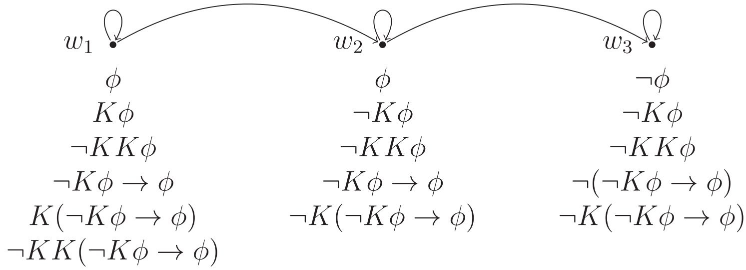

Fig. 1 KK failures and known abominable conditionals in non-transitive models

4. Iffy knowledge I: non-transitive models

This section investigates whether the following semantic analyses of the indicative conditional predict that failures of epistemic luminosity entail knowledge of abominable conditionals: the truth-functional analysis, the variably strict analysis, and the strict analysis.[12]

4.1 The truth-functional analysis

Let us begin with the truth-functional view, according to which a conditional of the form $\ulcorner \phi \to \psi \urcorner$ is true at *w* if and only if either ϕ is false at *w* or ψ is true at *w*.[13]

Material Conditional. $[\![\ulcorner \phi \to \psi \urcorner]\!]^w = 1$ if and only if $[\![\ulcorner \phi \supset \psi \urcorner]\!]^w = 1$.

This analysis serves as an easily verifiable test case for whether a given semantics for the indicative conditional yields knowledge of abominable conditionals on the present framework. Suppose that epistemic luminosity fails at some world in our non-transitive model. That is, $\ulcorner K\phi \wedge \neg KK\phi \urcorner$ is true at an arbitrary world *w*. Then, on the truth-functional view, the conditional $\ulcorner \neg K\phi \to \phi \urcorner$ is true at *w*, since both $\ulcorner K\phi \urcorner$ and ϕ are true there. Moreover, the conditional is true at any epistemically accessible world relative to *w*, since one's knowledge that

[12] We don't pretend this is a comprehensive survey of conditionals; a notable exception is the 'no-truth-value' approach popularized by Adams (1975), Edgington (1995), and Bennett (2003). One reason for the omission is that without a clear picture of how knowledge of conditionals works on this view, it's hard to test whether it secures (or fails to secure) knowledge of abominable conditionals.

[13] Proponents of this view include Lewis (1976), Jackson (1979, 1987), Grice (1989), Rieger (2006), and Williamson (2020).

ϕ at w ensures that ϕ is true at those worlds. Subsequently, the conditional is known at w, since it is true at every world epistemically accessible from w. Consequently, if the material conditional analysis of indicatives is correct, then relevant cases of abominable conditionals are known in non-transitive models.

4.2 Variably strict conditionals

Next, let us consider the view that indicative conditionals are *variably strict* conditionals. The central idea is that an indicative conditional is true at a world if and only if the consequent is true in all of the contextually determined set of worlds in which the antecedent is true. This restriction is specified by a (set of) contextual parameter(s) relative to which sentences are assigned semantic values, in addition to the usual world of evaluation. While these contextual parameters play no role in interpreting the non-iffy parts of the language, they do determine the set of worlds relevant to the evaluation of the consequent of the indicative conditional. Different implementations of this core idea constrain the contextually determined set of worlds in different ways. Whether failures of epistemic luminosity entail knowledge of abominable conditionals crucially depends on how the set of worlds is constrained. To illustrate this point, we shall consider Robert Stalnaker's and Angelika Kratzer's implementations of the view.

4.2.1 Stalnaker Stalnaker's analysis holds that a conditional of the form $\ulcorner \phi \rightarrow \psi \urcorner$ is true at a world w if and only if the closest ϕ-world to w is a ψ-world, where the closest ϕ-world to w is determined by a world-selection function f mapping pairs of sentences and worlds to worlds, subject to the following constraints:

(i) For all antecedents ϕ and base worlds w, ϕ must be true in $f(\phi, w)$.

(ii) For all antecedents ϕ and base worlds w, $f(\phi, w) = \lambda$ (the absurd world[14]) [if and] only if there is no world possible with respect to w in which ϕ is true.[15]

(iii) For all antecedents ϕ and base worlds w, if ϕ is true in w, then $f(\phi, w) = w$.

[14] The absurd world is the world in which everything is trivially true.

[15] See Mandelkern (2018) for arguments for this revision.

(iv) For all antecedents ϕ, ϕ' and base worlds w, if ϕ is true in $f(\phi', w)$ and ϕ' is true in $f(\phi, w)$, then $f(\phi, w) = f(\phi', w)$.

We can thus provide the following semantic clause for the indicative conditional (Stalnaker 1968, 1975; Stalnaker and Thomason 1970):

STALNAKER. $[\![\ulcorner\phi \rightarrow \psi\urcorner]\!]^{f,w} = 1$ if and only if $f(\phi, w) \in [\![\psi]\!]$.

Conditionals are also interpreted relative to the context set, roughly, the set of worlds characterizing each proposition taken for granted by all of the conversational participants. Since we only consider single-agent cases, for our purposes the context set at a world is determined by what one knows at that world: $C(w) = R(w)$.[16] On this basis, Stalnaker places a further pragmatic constraint on his theory to accommodate the difference between indicative and counterfactual conditionals, namely, that indicatives carry a default presupposition that every world in C is closer to every world in C than it is to any world outside C:

> If the conditional is being evaluated at a world in the context set, then the world selected must, if possible, be within the context set as well. ... In other words, all worlds within the context set are closer to each other than any worlds outside it. (Stalnaker 1975, pp. 275–6)

The formal upshot is that a selection function is admissible only if it satisfies the following condition:

PRAGMATIC CONSTRAINT. For all $w \in C$ and $z \notin C$, $f(\phi, w) = z$ only if $C \cap [\![\phi]\!] = \varnothing$.

Roughly speaking, a selection function picks out a world outside a context set only if the antecedent of the conditional is incompatible with the context set.

The third constraint above ensures that the world-selection function selects the world of evaluation, if the sentence is true in it: if $[\![\phi]\!]^{f,w} = 1$, then $f(\phi, w) = w$. The pragmatic constraint ensures that the conditional is defined only if $C \cap [\![\phi]\!] \neq \varnothing$. When the antecedent is not true at the base world, we must select the world that differs

[16] Knowledge-theoretic construals of the context set are a matter some controversy; see Hawthorne and Magidor (2009, 2010) and Stalnaker (2009). Nevertheless, we take the identification of the context set with the agent's knowledge to be less problematic in the present case, since we are concerned with whether agents, in and of themselves, know abominable conditionals.

minimally from the actual world. Our selection mandates that 'there are no differences between the actual world and the selected world except those that are required, implicitly or explicitly, by the antecedent' and that those changes do 'the least violence to the correct description and explanation of the actual world' (Stalnaker and Thomason 1970, p. 104).

With this idea of minimal change in mind, let's proceed to consider a failure of epistemic luminosity at an arbitrary world w in a non-transitive model. According to Stalnaker's semantics, the conditional $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at w, since (a) given that $\ulcorner \neg KK\phi \urcorner$ is true at w, there is some epistemically accessible world w' relative to w at which $\ulcorner \neg K\phi \urcorner$ is true, and (b) given that $\ulcorner K\phi \urcorner$ is true at w and w' is epistemically accessible relative to w, ϕ is true at w'. Thus ϕ is true at the closest $\ulcorner \neg K\phi \urcorner$-world relative to w.[17] The selection function does not pick worlds where $\ulcorner \neg K\phi \urcorner$ is true but ϕ is false, because such worlds are not (epistemically) accessible relative to w and they do not differ minimally from the actual world, since the truth-value of ϕ differs. The conditional is also true at any epistemically accessible world w' relative to w, since by constraint (ii), $f(\ulcorner \neg K\phi \urcorner, w') = w'$ and ϕ is true at w'. Consequently, the conditional is known at w, since it is true at every world epistemically accessible from it. Therefore, if the Stalnaker semantics for indicatives is correct, then abominable conditionals are known in non-transitive models.

4.2.2 Kratzer The dominant approach to the semantics of conditionals in linguistics is Kratzer's semantics (Kratzer 1981, 1986, 1991, 2012). According to Kratzer's account, a bare conditional like 'If Joey has a date tonight, then he is happy' is equivalent to 'If Joey has a date tonight, then it must be the case that he is happy'. Whether or not the modal *must* is covert, the function of the *if*-clause is to restrict the set of worlds over which *must* quantifies. The evaluation of the conditional is sensitive to two contextual parameters, which we call *conversational backgrounds*, functions from worlds to sets of propositions that uniquely determine the set of worlds over which the modal quantifies. For any world, the first conversational background, which we call the *modal base*, determines the set of worlds that are accessible

[17] The uniqueness condition fails if there is more than one epistemically accessible $\ulcorner \neg K\phi \urcorner$-world relative to w. In such cases, Stalnaker adopts a supervaluationist implementation of his semantics according to which a conditional is true at w if its consequent is true in all candidates for the closest antecedent world.

from that world, and the second conversational background, which we call the *ordering source*, induces a pre-order on the set of worlds accessible from that world. In general, a set of propositions A induces a pre-order $\leq_A$ on W as follows (Lewis 1981):

For all $w, w' \in W$, for any $A \subseteq \wp(W)$: $w \leq_A w'$ if and only if $\{p : p \in A \text{ and } w' \in p\} \subseteq \{p : p \in A \text{ and } w \in p\}$

That is, a world w is at least as close to another world w' relative to a set of propositions A if and only if all propositions of A which are true in w' are also true in w.

We can then provide the following semantic clause for indicative conditionals, relative to a modal base function f and an ordering source function g:[18]

KRATZER. $[\![\ulcorner \phi \rightarrow \psi \urcorner]\!]^{f,g,w} = 1$ if and only if $[\![\ulcorner \text{must } \psi \urcorner]\!]^{f',g,w} = 1$,

where $f'(w) = f(w) \cup \{[\![\phi]\!]^{f,g}\}$ and where $[\![\ulcorner \text{must } \phi \urcorner]\!]^{f,g,w} = 1$ if and only if for all $w' \in \cap f(w)$, if w' is minimal in $g_f(w)$ (the limitation of $g(w)$ on $\cap f(w)$), then $[\![\phi]\!]^{f,g,w'} = 1$.

These clauses are framed in terms too general to make concrete predictions about the present case, at least without further specifying the modal base and ordering source functions. Nevertheless, it is natural to suppose that the covert modal *must* is a strong epistemic modal, and so it should be evaluated against (i) an epistemic modal base (a function that maps a world to the set of propositions which one knows at that world), and (ii) either an empty ordering source function (the function that assigns the empty set to every world) or a normal ordering source function (the function that maps a world to the set of propositions that normally hold at that world) (cf. Kratzer 1986, 1991).[19] With these assumptions in place, we can show that when

[18] For simplicity, we make the limit assumption here, where Kratzer doesn't.

[19] An anonymous reviewer points out that the interpretation of *must* in KRATZER is a somewhat delicate matter. For example, it is well known that asserting a *must*-claim adds something to the proposition it embeds (cf. Karttunen 1972; von Fintel and Gillies 2010). However, Kratzer does not claim that an indicative conditional $\ulcorner \phi \rightarrow \psi \urcorner$ carries the same assertoric force, implicatures, and other general pragmatic effects as the assertion of $\ulcorner \text{must } \psi \urcorner$ on the supposition that ϕ, but rather that they are truth-conditionally equivalent. Nevertheless, it is worth investigating whether alternative accounts of *must* have different implications for our study. For example, we observe that if 'must' expresses something stronger than knowledge, such as knowledge of knowledge, then the abominable conditional may not even be true when one's knowledge that ϕ is not luminous. Contrastingly, if *must* is analysed against *similarity*-based ordering sources, then the abominable conditional may both be true and known. The verification of these results are left as an exercise for the reader.

one's knowledge that ϕ is not epistemically luminous at an arbitrary world *w*, the conditional $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at *w*.[20] However, even though the abominable conditional is true, one lacks knowledge of the abominable conditional, since for all one knows the conditional is false.[21] Consequently, $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is not known at *w*, because it is false at some epistemically possible world relative to *w*.

This result raises the question of why our investigation provides different results from Dorst's (2019 p. 1229).[22] For according to Dorst, a 'natural implementation' of Kratzer's analysis of the conditional is committed to Stability, and so, given the argument in §2, the conjunction of her analysis and Diligence should entail knowledge of abominable conditionals. However, the 'natural implementation' to which Dorst appeals invokes the following credal introspection principles in order to derive Stability from Kratzer's analysis:[23]

Positive Credal Introspection. $Cr(p) = 1 \supset Cr[Cr(p) = 1] = 1$

Negative Credal Introspection. $Cr(p) < 1 \supset Cr[Cr(p) < 1] = 1$

[20] According to Kratzer's semantics, the conditional $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at *w* if and only if ϕ is true in any $w' \in g_{f'}(w)$, the limitation of $g(w)$ on $\cap f'(w)$. Given that *f* is an epistemic modal base, it follows that $f'(w) = f(w) \cup \{[\![\ulcorner \neg K\phi \urcorner]\!]^{f,g}\} = \{p : S \text{ knows that } p \text{ in } w\} \cup \{[\![\ulcorner \neg K\phi \urcorner]\!]^{f,g}\}$. Given that *S* knows that ϕ in *w*, it follows that for every world $w' \in \cap f'(w)$, both ϕ and $\ulcorner \neg K\phi \urcorner$ are true at w'. It follows immediately that ϕ is true in every world in the limitation of $g(w)$ on $\cap f'(w)$, regardless of *g*.

[21] To show this, it is sufficient to identify some epistemically accessible world relative to *w* at which $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is not true. Recall that, since $\ulcorner \neg KK\phi \urcorner$ is true at *w*, there is some epistemically accessible world relative to *w* at which ϕ is true, but one does not know it. Call this world w', and let us reason about it directly. According to Kratzer's semantics, the conditional $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is false at w' if and only if ϕ is false at some world $w'' \in g_{f'}(w')$. Reasoning as before, given that *f* is an epistemic modal base, it follows that $f'(w) = \{p : \text{one knows that } p \text{ in } w\} \cup \{[\![\ulcorner \neg K\phi \urcorner]\!]^{f,g}\}$. But, unlike before, there is no guarantee that ϕ is true at every world in $\cap f'(w')$, since *S* does not know that ϕ in w'. Supposing that *g* is an empty ordering source, it follows immediately that not every world in the limitation of $g(w')$ on $\cap f'(w')$ will be a ϕ-world; hence the conditional will be false at w'. Furthermore, even if *g* is a 'normal' or 'stereotypical' ordering source (in the sense of Kratzer 1981), which would force us to consider what holds in the 'most normal' worlds in the epistemic modal base, it is plausible that there is some world in $g_{f'}(w')$ at which ϕ is false. After all, when we consider cases where one's knowledge fails to be luminous due to states of affairs in not-so-distant worlds, it is no stretch of the imagination to conclude that, *at such worlds*, the truth of ϕ is no more normal than its falsity. Once this is conceded, it follows immediately that the abominable conditional is false at some epistemically accessible world relative to *w*, and so one fails to know it.

[22] Thanks to an anonymous reviewer for encouraging us to clarify why our predictions differ from Dorst's.

[23] See Dorst (2019, p. 1243 n. 17) for proof.

Unlike Dorst's, our investigation does not assume the validity of any credal introspection principles. For when investigating the commitments of those who deny epistemic introspection principles, one should not immediately assume introspection principles of other kinds, such as credal introspection principles, in order to preserve the neutrality of the investigation. Indeed, those who reject epistemic introspection principles are likely to reject credal introspection principles for similar reasons, such as lacking the concept *credence* or *certainty*, or being in possession of only finitely many mental states. More worryingly, when combined with DILIGENCE, credal introspection principles mean that diligent agents enjoy non-trivial luminosity of their credences. For if one's credence in p is x, then the credal introspection principles guarantee that one is certain that one's credence in p is x; thus, by DILIGENCE, one *knows* that one's credence in p is x. Such luminosity is suspect in this context: it is dubious to argue for luminosity of one kind by assuming luminosity of another. Therefore, we do not assume these credal introspection principles in our models.

Without these principles, our model is inconsistent with STABILITY: given a diligent agent whose knowledge that ϕ is not luminous, it follows that $Cr(\phi) = 1$ and $Cr(\ulcorner \neg K\phi \urcorner) > 0$; but given that such an agent fails to know the corresponding abominable conditional $\ulcorner \neg K(\neg K\phi \rightarrow \phi) \urcorner$, it follows by DILIGENCE that $Cr(\ulcorner \neg K\phi \rightarrow \phi \urcorner) < 1$. This is a counterexample to STABILITY. For these reasons, we believe that proponents of Kratzer's analysis who reject introspection principles should reject STABILITY in full generality, rather than accept that abominable conditionals are knowable.

4.3 Strict conditionals

Lastly, consider the strict conditional account of the indicative. According to this view, an indicative conditional is true at a world w if and only if the corresponding material conditional is true at every R-accessible world relative to w:[24]

STRICT CONDITIONAL. $[\![\ulcorner \phi \rightarrow \psi \urcorner]\!]^{w} = 1$ if and only if for all w', if Rww', then $[\![\ulcorner \phi \supset \psi \urcorner]\!]^{w'} = 1$.

[24] Our inclusion of the strict conditional analysis is guided by Dorst's consideration of the view. Note that Kratzer's semantics is equivalent to a strict conditional analysis if both the modal base and ordering source are empty (Kratzer 1991, p. 649)

We must make some stipulations about the accessibility relation R relevant for the evaluation of conditionals. Suppose the accessibility relation is epistemic. Then suppose w is an arbitrary world in which one's knowledge that ϕ is not luminous. To see whether $\ulcorner\neg K\phi \rightarrow \phi\urcorner$ is known at w, we must see whether $\ulcorner\neg K\phi \rightarrow \phi\urcorner$ is true at every epistemically accessible world relative to w. Given our results for the material conditional, we already know $\ulcorner\neg K\phi \supset \phi\urcorner$ is true at every epistemically accessible world relative to w, and so the conditional is true at w. But given that $\ulcorner\neg KK\phi\urcorner$ is true at w, there is some epistemically accessible world w' relative to w at which $\ulcorner\phi \wedge \neg K\phi\urcorner$ is true. The conditional $\ulcorner\neg K\phi \rightarrow \phi\urcorner$ is not true at w', since there is some epistemically accessible world w'' relative to w' such that (i) ϕ is false at w'', since $\ulcorner\neg K\phi\urcorner$ is true at w', and so (ii) $\ulcorner\neg K\phi\urcorner$ is true at w''. That is, there is some epistemically accessible world relative to w' at which $\ulcorner\neg K\phi \supset \phi\urcorner$ is false. Thus the indicative conditional is false at w', and so it is not known at w. Therefore, if the strict conditional analysis of indicatives is correct, then there is no knowledge of abominable conditionals in non-transitive models.[25]

An anonymous reviewer points out that anyone who accepts the following two principles must also accept KK:[26]

Strict. Necessarily, $\ulcorner\phi \rightarrow \psi\urcorner$ if and only if $\ulcorner K(\phi \supset \psi)\urcorner$

Restricted Or-To-If (ROTI). Necessarily, if $\ulcorner\neg K\neg\phi\urcorner$ and $\ulcorner K(\phi \supset \psi)\urcorner$, then $\ulcorner K(\phi \rightarrow \psi)\urcorner$.

This is worth noting because, as the reviewer points out, Strict is just the strict analysis of the indicative conditional (interpreted with an epistemic accessibility relation), and anyone who accepts Diligence and Stability will be committed to ROTI. Consequently, proponents of the strict analysis must reject ROTI, and so also reject either Diligence or Stability, in order to avoid epistemic luminosity.[27]

[25] Interestingly, the material conditional is a limiting case of the strict conditional, where the accessibility relation R requires worlds only to see themselves.

[26] *Proof.* Let ϕ be any tautology and suppose $\ulcorner K\psi\urcorner$. Assume K has a normal modal logic. Since ϕ is a tautology, it follows that $\ulcorner K\phi\urcorner$, and so $\ulcorner\neg K\neg\phi\urcorner$. And since $\ulcorner K\psi\urcorner$, it trivially follows that $\ulcorner K(\phi \supset \psi)\urcorner$. Then by ROTI, $\ulcorner K(\phi \rightarrow \psi)\urcorner$. By the K axiom and *modus ponens*, $\ulcorner K\phi \rightarrow K\psi\urcorner$. And by Strict, $\ulcorner K(K\phi \supset K\psi)\urcorner$. By another instance of the K axiom and *modus ponens*, $\ulcorner KK\phi \supset KK\psi\urcorner$. And since ϕ is a tautology, $\ulcorner KK\phi\urcorner$, and thus $\ulcorner KK\psi\urcorner$.

[27] For further discussion of the connection between epistemic luminosity and ROTI, see, among others, Rothschild and Spectre (2018) and Holguín (2021).

4.4 A summary of results

In this section, we have surveyed the major theories of conditionals, namely, the truth-functional analysis, variably strict analyses, and the strict conditional analysis. Interestingly, knowledge of abominable conditionals is secured only on the truth-functional analysis and Stalnaker's analysis, but false on Kratzer's analysis and on strict conditional semantics, as summarized in Table 1, where w is a world at which one's knowledge that ϕ is not luminous.

Of course, not all lines in our table are created equal. For example, one should perhaps not read too much into the fact that the truth-functional analysis entails knowledge of abominable conditionals, especially given its waning influence in recent years. Nevertheless, it is worth noting that more prominent analyses, such as the Stalnaker and Kratzer accounts, differ on whether they entail knowledge of abominable conditionals. Consequently, epistemic non-luminosity by no means guarantees knowledge of abominable conditionals.

5. Iffy knowledge II: transitive models

Knowledge of abominable conditionals is by no means guaranteed for epistemically non-luminous agents. But how does it fare for epistemically luminous agents? Our framework provides an ideal setting to investigate this question, since epistemic luminosity can be modelled in epistemic logic by any Kripke model in which the epistemic accessibility relation is at least both reflexive and transitive.[28] The smallest such Kripke models are S4 models. Using these models, this section investigates whether epistemic luminosity avoids knowledge of abominable conditionals by surveying the theories outlined in the previous section.

It turns out that on transitive models, if one knows that ϕ, then one knows the abominable conditional ⌜If I don't know that ϕ, then ϕ⌝. The reason for this is that all theories of the conditional that we consider validate the following condition:[29]

Vacuity. For any antecedent ϕ and consequent ψ, if $R(w) \cap [\![\phi]\!] = \varnothing$ then ⌜$\phi \rightarrow \psi$⌝ is true at w.[30]

[28] The epistemic accessibility relation is transitive if and only if $\forall xyz(Rxy \wedge Ryz \supset Rxz)$.

[29] Thanks to an anonymous reviewer for suggesting this formulation of our point.

[30] *Proof.* (1) *Material conditional.* On the material conditional analysis, $R(w) = \{w\}$. Therefore if ϕ is impossible at w, ϕ is false at w and so by the truth conditions for the

Table 1 Summary of conditionals in non-transitive models

Semantic theory	Truth-value of $\ulcorner K(\neg K\phi \rightarrow \phi)\urcorner$ in w
Material conditional	T
Stalnaker	T
Kratzer	F
Strict conditional	F

In other words, if ϕ is epistemically impossible at your world, then $\ulcorner \phi \rightarrow \psi \urcorner$ is true at your world, regardless of the content of ψ.

VACUITY, transitive models and knowledge of some proposition suffice for knowledge of the corresponding abominable conditional. Suppose $\ulcorner K\phi \urcorner$ is true at *w*. By transitivity, $\ulcorner K\phi \urcorner$ is true at every world accessible from *w*. Therefore $\ulcorner \neg K\phi \urcorner$ is impossible at *w*, and so $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at *w*. Because the frame is transitive, the same reasoning holds for every world accessed by *w*. Since $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at every $w \in R(w)$, $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is known at *w*, on every analysis of the indicative we have considered. These results are summarized in Table 2.

Importantly, our argument has not relied on any potentially controversial bridging principles between knowledge and credence. Rather, these results fall out of standard transitive models for epistemic logic and certain theories of the indicative. The upshot is that proponents of epistemic luminosity must appeal to extra-semantic considerations if they wish to avoid knowledge of abominable conditionals.[31]

material conditional, $\ulcorner \phi \rightarrow \psi \urcorner$ is true at *w*. (2) *Stalnaker analysis*. By constraint (ii) (see §4.2.1), if the antecedent ϕ is impossible, then the selection function selects the absurd world λ: $f([\![\phi]\!], w) = \lambda$. Since by definition everything is true in the absurd world, the consequent is true at the absurd world and so the conditional is vacuously true at *w*. (3) *Kratzer*. On this analysis, the conditional is vacuously true at *w* if there are no accessible worlds compatible with the antecedent. (4) *Strict conditional*. On this view, $\ulcorner \phi \rightarrow \psi \urcorner$ is true just in case $\ulcorner \phi \supset \psi \urcorner$ is true at all $w \in R(w)$. Since ϕ is impossible at *w*, then ϕ is false at every $w \in R(w)$. Therefore $\ulcorner \phi \supset \psi \urcorner$ is true at every $w \in R(w)$. Therefore $\ulcorner \phi \rightarrow \psi \urcorner$ is true at *w*.

[31] An anonymous reviewer suggests that VACUITY may provide the proponent of epistemic luminosity a privileged route to explaining the infelicity of abominable conditionals. For, as the reviewer observes, when one's knowledge that *p* is epistemically luminous, VACUITY predicts that one will also know the uncontroversially bizarre-sounding 'If I don't believe that *p*, then I know that *p*' and 'If I don't know that *p*, then I know that *p*'. The reviewer notes that proponents of epistemic luminosity can take heart from the fact that whatever explanation we provide to these conditionals to explain their badness, the same explanation can be carried

Table 2 Summary of conditionals in transitive models

Semantic theory	Truth-value of $\ulcorner K(\neg K\phi \rightarrow \phi)\urcorner$ at w
Material conditional	T
Stalnaker	T/U
Kratzer	T
Strict conditional	T

We envision two objections to the argument just given. The first objection complains that analyses of the indicative should not validate VACUITY; that by appealing to this feature of the analyses, we are somehow doing something illegitimate. There are several things to say in reply. First, our goal in this paper is simply to combine epistemic logic and standard semantic analyses of the conditional in a straightforward way, to see what we can learn about the knowability of abominable conditionals. Our resources have deliberately been austere: if one wants to reject this VACUITY feature of most conditional analyses, one does so at the cost of serious violence to their semantics. Now, one might regard this violence as an acceptable cost of blocking VACUITY. This raises our second concern with this objection. If indicatives with impossible antecedents are not to be vacuously true at *w*, then there are two prima facie plausible options: either those conditionals are false or they are truth-valueless. The problem with either of these proposals is that they predict that, for any (epistemically) impossible ϕ, $\ulcorner \phi \rightarrow \phi \urcorner$ is either false or true-valueless. But such a conditional should be trivially true, something VACUITY makes ready sense of. Of course, there may be more sophisticated proposals for removing VACUITY. But the worry just

over to abominable conditionals. They further note that this dialectical move is not available to opponents of epistemic luminosity, at least not given theories of the conditional which make such conditionals non-vacuously true and known. If this is correct, then proponents of epistemic luminosity are in a uniquely good position to explain why abominable conditionals are abominable.

However, we believe that non-uniform explanations should be given for the badness of abominable conditionals and the conditionals above. In particular, we would like to suggest that the source of the infelicity of the conditionals above stems from the fact that the antecedents and consequents are each independently semantically consistent, but jointly semantically inconsistent. For example, one cannot both know that not *p* and know that *p*, one cannot both not believe that *p* and know that *p*, and so on. It is this semantic inconsistency that our (in)felicity judgements are tracking in this case. If something along these lines is correct, then this explanation can be exploited regardless of whether knowledge is luminous: anyone can explain the badness of these conditionals, and the explanation does not depend on VACUITY.

mentioned shows the dangers of tampering with well-behaved and plausible semantic accounts of the indicative merely to block some apparently unrelated conclusions regarding the correct logic for knowledge. Such vandalism is unlikely to be cost-free.

The second objection tells an alternative, more demanding story about what is required for *knowledge* of abominable conditionals; a story that interferes with our conclusion. Rather than alter the truth conditions of the indicative, one might argue that one's knowledge of a conditional should be *robust* in the sense that, if one's knowledge is minimally weakened to leave open the antecedent, one must still know the conditional. A natural formalization of this idea is given as follows (cf. Dorst 2019, p. 1236):

ROBUSTNESS. $\ulcorner K(\phi \rightarrow \psi)\urcorner$ if and only if $\ulcorner K_\phi(\phi \rightarrow \psi)\urcorner$, where $\ulcorner K_\phi \urcorner$ is one's information state minimally revised to include the closest ϕ-possibilities.

Intuitively, ROBUSTNESS holds that one can know a conditional $\ulcorner \phi \rightarrow \psi \urcorner$ if and only if one knows that conditional when one's epistemically accessible worlds are expanded to include the nearest ϕ worlds. So when one has an epistemically impossible antecedent, one only gets to know the conditional if, upon adding the nearest antecedent worlds to one's set of epistemically accessible worlds, one still knows it.

Dorst (2019, pp. 1237–8) shows that, given transitive models like the ones we have considered, ROBUSTNESS suffices for the unknowability of abominable conditionals. In other words:

For any ϕ and w, if $[\![\phi]\!] \neq W$, then $[\![\ulcorner \neg K(\neg K\phi \rightarrow \phi)\urcorner]\!]^w = 1$

That they cannot be known provides a straightforward explanation of their infelicity: asserting them violates the knowledge norm of assertion. However ROBUSTNESS is considerably more costly than has been realized. First, as we have in effect seen, ROBUSTNESS is inconsistent with the standard semantics for K and the various semantics for indicatives we have considered; those two things jointly predict that abominable conditionals are known on transitive models.

A more dialectically neutral point is that given a semantics for the indicative that validates STRONG CENTRING and VACUITY, such as the Stalnaker analysis, ROBUSTNESS gives rise to a kind of sceptical consequence:

STRONG CENTRING. For any antecedent ϕ, if $w \in [\![\phi]\!]$, then for all $w' \neq w$, $w <_{g(w)} w'$.

It turns out that, given these constraints on the indicative semantics, $[\![\phi]\!] = [\![\ulcorner \neg K\phi \rightarrow \phi \urcorner]\!]$; in other words, the two sentences have identical truth conditions.[32]

Given Robustness and a transitive model, we have Dorst's result that for any non-tautologous ϕ, $\ulcorner \neg K(\neg K\phi \rightarrow \phi)\urcorner$. By our earlier reasoning, $\ulcorner \phi \supset (\neg K\phi \rightarrow \phi)\urcorner$ is a theorem. By the necessitation axiom in normal modal logics, $\ulcorner K(\phi \supset (\neg K\phi \rightarrow \phi))\urcorner$. An application of the K axiom and *modus ponens* yields $\ulcorner K\phi \supset K(\neg K\phi \rightarrow \phi)\urcorner$. Dorst's result and *modus tollens* yields $\ulcorner \neg K\phi \urcorner$. That consequence is global scepticism.

So if Stalnaker's analysis is correct, the knowability of the abominable conditionals cannot be blocked by appeal to Robustness, on pain of global scepticism. This result is bad for the defender of epistemic luminosity for two reasons. The first is that it means they are apparently committed to knowledge of abominable conditionals on virtually all popular analyses of the indicative. The second is that if the defender of epistemic luminosity wants to hold on to Robustness and thereby endorses Kratzer's or the strict conditional analysis, then they are moving to precisely those theories of the conditional which do not give rise to knowledge of abominable conditionals on the non-transitive models with which we started out. In either case, the proponent of epistemic luminosity is firmly on the back foot.

6. Know it, but don't say it

Regardless of whether epistemic luminosity fails, some theories of the indicative fail to secure knowledge of abominable conditionals. This provides an easy explanation of the unassertability of abominable conditionals: given the knowledge norm of assertion, if one doesn't

[32] *Proof.*

$\Rightarrow$ Suppose $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at w (for example, $w \in [\![\ulcorner \neg K\phi \rightarrow \phi \urcorner]\!]$). Since indicatives entail the corresponding material conditional, $\ulcorner K\phi \vee \phi \urcorner$ is true at w and therefore ϕ is true at w (that is, $w \in [\![\phi]\!]$).

$\Leftarrow$ Suppose Strong Centring holds, and that ϕ is true at w (for example, $w \in [\![\phi]\!]$). One of $\ulcorner K\phi \urcorner$ and $\ulcorner \neg K\phi \urcorner$ is true at w. If $\ulcorner \neg K\phi \urcorner$ is true at w, then by Strong Centring, $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at w. If $\ulcorner K\phi \urcorner$ is true at w then either there is some accessible $\ulcorner \neg K\phi \urcorner$-world or there is not. If the former, then that accessible $\ulcorner \neg K\phi \urcorner$-world is guaranteed to be a ϕ-world, by the fact that $\ulcorner K\phi \urcorner$ is true at w. In that case, $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at w, since the nearest accessible $\ulcorner \neg K\phi \urcorner$-world is a ϕ-world. Alternatively, there is no accessible $\ulcorner \neg K\phi \urcorner$-world, in which case $\ulcorner \neg K\phi \rightarrow \phi \urcorner$ is true at w by Vacuity. In either case, $w \in [\![\ulcorner \neg K\phi \rightarrow \phi \urcorner]\!]$.

know an abominable conditional, then it is impermissible to assert it. However, knowledge of abominable conditionals turns out to plague both houses of epistemic luminosity: regardless of whether one denies epistemic luminosity, there are some theories of the indicative conditional that give rise to knowledge of abominable conditionals. These results rob us of an easy explanation for the unassertability of abominable conditionals. In this section, we explore whether there are other ways to explain away the infelicity of abominable conditionals.

Assertions can be infelicitous even if their contents are known, since knowledge is only a *necessary* condition on assertion. For example, absent a degenerate context, sentences like (2) are plainly infelicitous, even if one can know them:[33]

(2) #If there's life on Mars, I'm sitting at my desk.

Abominable conditionals aside, then, theorists of the conditional must explain why such conditionals are infelicitous. It is natural to hope that such explanations can be assimilated to abominable conditionals.

The standard explanation of the infelicity of (2) is that the antecedent lacks any relevance to the consequent, and when the antecedent of a conditional lacks relevance to its consequent in the context, the conditional is likely to be rejected. Something similar arguably holds of the abominable conditional: *the supposition that one does not know that* ϕ *is not generally relevant to whether* ϕ *is true*. That is, assertions of abominable conditionals generally implicate that the speaker's failure to know the consequent is somehow relevant to the truth of the consequent. But, given our understanding about mind–world directions of fit, this implication is generally implausible: other things being equal, that we don't know p is irrelevant to whether p is true. Consequently, unless there is a reason to think that one's lack of knowledge about some proposition is relevant to the truth of that proposition, the relevance implicature carried by standard abominable conditionals strikes us as implausible, and so such conditionals are infelicitous.

While a systematic development of the notion of relevance is outside the scope of this paper, this explanation is supported by the fact that one's epistemic position can be relevant to the salient features of reality, at least insofar as it sometimes provides evidence for what

[33] Indeed, given knowledge of the consequent and non-zero credence in the antecedent, something like Stability entails that diligent agents know (2).

those features are. Consider, for example, sentences like (3) or its less verbose equivalent, (4):

(3) If I don't know there's a proposition that I don't know, then there's a proposition I don't know.[34]

(4) If I don't know I am ignorant of something, then I am ignorant of something.

That these conditionals are both knowable and assertable, despite that fact they have the form $\ulcorner \neg K\phi \rightarrow \phi \urcorner$, is testament to the fact that not every such conditional is 'abominable', and so militates against any theory that predicts as much. Importantly, considerations of relevance do not rule these sentences as unassertable, since one's ignorance about what one knows counts as evidence for the fact that one is ignorant in the first place, and so the antecedent is relevant to the consequent.[35]

A natural worry for this approach is that it cannot handle variants of abominable conditionals like the following (cf. Dorst 2019, p. 1231):

(5) #Even if I don't know that Padua is in Italy, Padua is in Italy.

(6) #Whether or not I know that Padua is in Italy, Padua is in Italy.

Assuming that the meanings of these conditionals are minimally different from normal indicative conditionals, it follows from our arguments above that agents may sometimes know these variants too. But since *even if* and *whether or not* modifiers cancel any expectation of relevance, the explanation for infelicity of abominable conditionals does not generalize to these variants (cf. Dorst 2019, p. 1231).

While we agree that these modifiers have this relevance-cancelling function, we disagree that a uniform treatment of their infelicity is required. For in addition to cancelling relevance, one can assert

[34] Thanks to John Hawthorne for bringing this example to our attention.

[35] Conditionals like (3) and (4) are also assertable even in contexts that explicitly disavow any presupposition of the consequent, such as (i):

(i) Who knows whether Joey is ignorant, but if he doesn't know he's ignorant, he's ignorant.

Consequently, the felicity of these conditionals does not hinge on any presupposition of the consequent in the context; cf. Dorst (2019, p. 1231).

them only if, roughly speaking, one knows the consequent.[36] Compare:

(7) I don't know whether Monica is going to the party but I do know that if Ross goes, she'll go.

(8) #I don't know whether Monica is going to the party but I do know that even if Ross goes, she'll go.

(9) #I don't know whether Monica is going to the party but I do know that whether or not Ross goes, she'll go.

Given that one can only assert *even if* and *whether or not* conditionals in cases where one knows their consequent, there is a straightforward explanation of the infelicity of these variants of abominable conditionals: in asserting them, one presupposes a possibility that one simultaneously implicitly undermines, namely, that one knows the consequent.[37] Consequently, these variants convey something similar to Moorean-esque sentences like ⌜ϕ but I might not know that ϕ⌝, which are arguably infelcitous because the latter conjunct 'blatantly undermines the normally intended effect of the first conjunct on the audience, by giving reason not to rely on it' (Williamson 2013, p. 82). Contra Dorst (2019, p. 1253), a parallel explanation is available for the infelicity of the *even if* and *whether or not* variants of the abominable conditional, though not Dorst's original example, since an assertion of the latter does not presuppose that the speaker knows the consequent.

It is a general feature of *even if* and *whether or not* conditionals that their warranted assertion requires knowledge of their consequents; neither the antecedent nor consequent needs to concern one's knowledge. Compare:

(10) #If Joey doesn't have a date, then Joey's date will be sad.

The natural explanation of the infelicity of (10) is that one presupposes the possibility that Joey does not have a date while simultaneously disavowing that possibility. Given this general feature of *even if* and *whether or not* conditionals, there is no need to assimilate an explanation of their infelicity to our explanation of the infelicity of standard abominable conditionals.

[36] See Guerzoni and Lim (2007) for good discussion of this property of *even if* conditionals.

[37] Dorr and Hawthorne (2013, pp. 888–9) make a similar point, that 'in ordinary settings, there is something quite odd about avowing that one fails to know whether one knows'.

One may also worry that our initial explanation cannot explain the infelicity of third-personal abominable conditionals like the following:[38]

(11) #If Joey doesn't know it's Wednesday, it's Wednesday.

But the explanation of this infelicity is again just a lack of relevance between antecedent and consequent. When these conditionals are transposed to *even if* or *whether or not* conditionals, settings that disavow this sort of connection, the infelicity disappears.

(12) Even if Joey doesn't know it's Wednesday, it's Wednesday.

(13) Whether or not Joey knows it's Wednesday, it's Wednesday.

The upshot is that considerations of relevance explain the infelicity of abominable conditionals, even if one's favoured theory of the indicative allows their knowledge.

Finally, one may worry that even if we can explain the infelicity of abominable conditionals, we have said little about how our discussion generalizes to other cognitive attitudes. For example, in so far as certain theories of conditionals and assumptions about the structure of knowledge conspire to generate knowledge of abominable conditionals, they also conspire to generate belief in abominable conditionals, which would be just as bad.

Let us conclude, then, by noting some general constraints on how these facts should be explained. First, it is worth considering an analogy with conditionals like (2), which are believable when one knows the consequent and something like Stability holds. In our experience, the studious sedentary person is likely to be perplexed if asked whether they believe that if there's life on Mars, they're sitting at their desk: most people do not take themselves to have such beliefs. It may be that we erroneously misjudge that we do not believe such conditionals when in fact we do. Or the prediction that we believe them may be noise generated from idealizing assumptions built into conditions like Stability. In any case, given the plausibility of assimilating relevance-based explanations for the infelicity of (2) to abominable conditionals, we require special reason for thinking that whatever explains why (2) seems unbelievable cannot also be extended to the apparent strangeness of believing abominable conditionals.

[38] Thanks to an anonymous reviewer for raising this point.

Second, blanket prohibitions on believing abominable conditionals seem too strong. For we saw above that any theory that validates VACUITY and STRONG CENTRING, such as Stalnaker's theory, equates the truth conditions of ϕ and $\ulcorner \neg K\phi \rightarrow \phi \urcorner$. Therefore, given weak closure conditions on belief, if one believes ϕ, then one also believes $\ulcorner \neg K\phi \rightarrow \phi \urcorner$. Consequently, any blanket ban on believing abominable conditionals entails total agnosticism: agents would be maximally unopinionated about any non-trivial proposition. Moreover, any blanket ban on believing abominable conditionals cuts against the good standing of (3) and (4), conditionals that are felicitous, knowable, believable, and perhaps deserving of the status of theoremhood, despite being of 'abominable' form.

Lastly, those who deny epistemic luminosity should not be so worried about believing abominable conditionals. After all, failures of epistemic luminosity often entail that we may justifiably believe strange conjunctions of the form $\ulcorner \phi$ but I'm not sure I should believe $\phi \urcorner$. Such beliefs are explained by rejecting any general anti-akratic requirement that first-order beliefs should cohere with higher-order attitudes about what one ought to believe.[39] In general, such theorists will be unperturbed by incurring additional strange beliefs from failures of epistemic luminosity. While these consequences must be explained, the presence of some intuitively strange beliefs is priced into their position. This is decidedly not the case for proponents of epistemic luminosity, one of the motivations for which was supposed to be that it can block these weird beliefs from arising in the first place.

7. Concluding remarks

This paper has argued that whether epistemic luminosity or epistemic non-luminosity entails knowledge of abominable conditionals is highly sensitive to the semantic analysis of the indicative conditionals under discussion. Whether knowledge is luminous essentially restricts the available space of correct semantic analyses for the indicative conditional. In this sense, both advocates of epistemic luminosity and

[39] For developments of this kind of 'level-splitting' view, see, among others, Lasonen-Aarnio (2014) and Weatherson (2019). Broadly similar remarks apply, we think, to some of the weird patterns of reasoning apparently licensed by abominable conditionals; cf. Dorst (2019, p. 1230). Level-splitters already have to say something about the strange patterns of reasoning that arise in cases of luminosity failure; see Weatherson (2019, ch. 9) for good discussion.

dissenters from it can avoid knowledge of abominable conditionals with careful selection of their semantics. We have also suggested that standard pragmatic resources can explain the infelicity even if the abominable conditionals can be known. In sum, the iffy argument for epistemic luminosity may be just that: iffy.[40]

References

Adams, Ernest W. 1975: *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Dordrecht: D. Reidel.

Bennett, Jonathan 2003: *A Philosophical Guide to Conditionals*. Oxford: Oxford University Press.

Das, Nilanjan, and Bernhard Salow 2018: 'Transparency and the KK Principle'. *Noûs*, 52(1), pp. 3–23.

Dorr, Cian, and John Hawthorne 2013: 'Embedding Epistemic Modals'. *Mind*, 122, pp. 867–913.

Dorst, Kevin 2019: 'Abominable KK Failures'. *Mind*, 128, pp. 1227–59.

Edgington, Dorothy 1995: 'On Conditionals'. *Mind*, 104, pp. 235–329.

Feldman, Richard 1981: 'Fallibilism and Knowing that One Knows'. *Philosophical Review*, 90(2), pp. 266–82.

von Fintel, Kai, and Anthony S. Gillies 2010: '*Must* … Stay … Strong!'. *Natural Language Semantics*, 18(4), pp. 351–83.

Goodman, Jeremy, and Bernhard Salow 2018: 'Taking a Chance on KK'. *Philosophical Studies*, 175(1), pp. 183–96.

Greco, Daniel 2014: 'Could KK Be OK?' *Journal of Philosophy*, 111(4), pp. 169–97.

—— 2015: 'Iteration and Fragmentation'. *Philosophy and Phenomenological Research*, 91(3), pp. 656–73.

Grice, Paul 1989: *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Guerzoni, Elena, and Dongsik Lim 2007: '*Even if*, Factivity and Focus'. In E. Puig-Waldmüller (ed.), *Proceedings of Sinn und Bedeutung* 11, pp, 276–90. Barcelona: Universitat Pompeu Fabra.

Hawthorne, John, and Ofra Magidor 2009: 'Assertion, Context, and Epistemic Accessibility'. *Mind*, 118, pp. 377–97.

[40] We should like to thank John Hawthorne, Lukas Lewerentz, Matthew Mandelkern, Weng Kin San and Timothy Williamson for extremely helpful comments on a previous draft of this paper, with special thanks to Richard Roth for numerous valuable discussions and comments. Matthew Hewson would also like to acknowledge financial support from the Clarendon Fund and Merton College. The authors contributed equally to the ideas throughout this paper; authorship is listed in alphabetical order.

—— 2010: 'Assertion and Epistemic Opacity'. *Mind*, 119, pp. 1087–1105.
Hintikka, Jaakko 1962: *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press.
Holguín, Ben 2021: 'Indicative Conditionals without Iterative Epistemology'. *Noûs*, 55(3), pp. 560–80.
Jackson, Frank 1979: 'On Assertion and Indicative Conditionals'. *Philosophical Review*, 88(4), pp. 565–89.
—— 1987: *Conditionals*. Oxford: Basil Blackwell.
Karttunen, Lauri 1972: 'Possible and Must'. In J. P. Kimball (ed.), *Syntax and Semantics*, Volume 1, pp. 1–20. New York: Academic Press.
Kratzer, Angelika 1981: 'The Notional Category of Modality'. In Hans J. Eikmeyer and Hannes Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pp. 38–74. Berlin: De Gruyter.
—— 1986: 'Conditionals'. *Proceedings of the Chicago Linguistics Society*, 22, pp. 1–15.
—— 1991: 'Modality'. In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, pp. 639–50. Berlin: De Gruyter Mouton.
—— 2012: *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
Lasonen-Aarnio, Maria 2014: 'Higher-Order Evidence and the Limits of Defeat'. *Philosophy and Phenomenological Research*, 88(2), pp. 314–45.
Lewis, David 1976: 'Probabilities of Conditionals and Conditional Probabilities'. *Philosophical Review*, 85(3), pp. 297–315.
—— 1981: 'Ordering Semantics and Premise Semantics for Counterfactuals'. *Journal of Philosophical Logic*, 10(2), pp. 217–34.
Mandelkern, Matthew 2018: 'The Case of the Missing "If": Accessibility Relations in Stalnaker's Theory of Conditionals'. *Semantics and Pragmatics*, 11, pp. 1–5.
Rieger, Adam 2006: 'A Simple Theory of Conditionals'. *Analysis*, 66(3), pp. 233–40.
Rothschild, Daniel, and Levi Spectre 2018: 'A Puzzle about Knowing Conditionals'. *Noûs*, 52(2), pp. 473–78.
Stalnaker, Robert C. 1968: 'A Theory of Conditionals'. In Nicholas Rescher (ed.), *Studies in Logical Theory*, pp. 98–112. Oxford: Oxford University Press.
—— 1975: 'Indicative Conditionals'. *Philosophia*, 5(3), pp. 269–86.

—— 1999: *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.
—— 2009: 'On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility'. *Mind*, 118, pp. 399–409.
Stalnaker, Robert C. and Richmond H. Thomason 1970: 'A Semantic Analysis of Conditional Logic'. *Theoria*, 36(1), pp. 23–42.
Weatherson, Brian 2019: *Normative Externalism*. Oxford: Oxford University Press.
Williamson, Timothy 1992: 'Inexact Knowledge'. *Mind*, 101, pp. 217–42.
—— 2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.
—— 2013: 'Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic'. *Inquiry*, 56(1), pp. 77–96.
—— 2014: 'Very Improbable Knowing'. *Erkenntnis*, 79(5), pp. 971–99.
—— 2020: *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford: Oxford University Press.

