Human Foreknowledge

Fabrizio Cariani
University of Maryland, College Park
forthcoming in *Philosophical Perspectives*

Abstract

I explore the motivation and logical consequences of the idea that we have some (limited) ability to know contingent facts about the future, even in presence of the assumption that the future is objectively unsettled or indeterminate. I start by formally characterizing skepticism about the future. This analysis nudges the anti-skeptic towards the idea that if some propositions about the future are objectively indeterminate, then it may be indeterminate whether a suitably positioned agent knows them.

1 Introduction

Here are some things I know about my own future. Of a crystal glass that is dropped from the tenth floor, I know that it will shatter. I know that it will snow in Chicago next winter; that I will drink some water within the next twelve hours; that neither my spouse nor I will give birth next month; that Paris will be in France a year from now. I also know many things about the future that are bolder and more controversial than these, but will not list them here so as to not distract

[0]For conversation and exchanges, I am grateful to Mike Caie, Ivano Ciardelli, Bill D'Alessandro, Simon Goldstein, Eric Pacuit, Paolo Santorio, Rachel Singpurwalla, Allen Stairs, Stephan Torre, and Robbie Williams. I thank audiences at the UMD Work in Progress series, as well as the Logic Colloquium at the Munich Center for Mathematical Philosophy. The paper was also improved by online exchanges with Andrew Bacon, Rohan French, and Jack Woods. A Facebook thread (!) I started in the group *Board Certified Epistemologists* helped me see some pitfalls in some claims I had been tempted to make. I was especially helped by comments and pointers by Sam Carter, Georgi Gardiner, Clayton Littlejohn, Ángel Pinillos, Julia Staffel, Ram Neta, and Alex Worsnip. Paolo Santorio, Simon Goldstein and Eric Pacuit get a second dose of thanks: Paolo and Simon read and commented on an early draft of the paper; Eric helped with pointers to literature on fusion and with a proof. Finally, Brian Rabern read a later version of the manuscript and corrected several important oversights.

you from my main point. If knowledge has an important role to play within our understanding of assertion, action, thought, and inquiry, we need to be able to make sense of the possibility for a non-divine, non-magical intellect to know the future.

The countervailing thought is that fundamental asymmetries between past and future might prevent us from knowing contingent propositions that are entirely about the future (henceforth *future contingents*). Indeed, the idea that humans can never know contingent propositions about the future is an intellectual trope, even independently of any philosophical commitment. Call the position that rejects knowledge of future contingents *future skepticism*. Global skeptics—those who deny that we have any knowledge—are, *inter alia*, future skeptics. However, the more interesting kind of future skeptic to be targeted here holds that we have copious amounts of knowledge about present and past facts, but lack knowledge of the future altogether.

What kinds of consideration might warrant future skepticism? It is often mentioned that our knowledge of the future is limited because of the asymmetry of causation—that is, because causes do not seem to proceed from later to eventualities to earlier ones. However, while severely constraining of our ability to acquire empirical knowledge of the future, the asymmetry of causation does not justify future skepticism, except for those who are willing to start from extreme empiricist premises.[1] Even Goldman — a philosopher with strongly naturalistic views — makes specific room for human foreknowledge by amending his causal theory of knowledge specifically to allow beliefs about the future to qualify as knowledge (Goldman, 1967, pp.364-365). For Goldman, Anna is in a position to know that the ice cream will melt if there is a common cause (of the right kind) of both Anna's belief that the ice cream will melt and of some eventuality that settles it as true. Setting the specific suggestion aside, the general point is that the asymmetry of causation is a strong constraint on foreknowledge, but not an impassable barrier.

More threatening challenges emerge when we reflect on the design of an interface between a theory of knowledge and a theory of objective chance. It is widely accepted that knowledge requires truth. It might, however, require more than that. For example, it might be impossible to know propositions that that have

[1] As I argue in Cariani (2021, ch.12), it is valuable to ask after a naturalistic explanation of how foreknowledge might be acquired. But we can see the general outlines of how such an explanation might go. We can know future facts in part by inductive inference, and in part by some kind of constrained imagined activity that has been independently crucial to counterfactual cognition (see Kahneman and Tversky, 1982; Williamson, 2008; Balcerak Jackson, 2018; Aronowitz and Lombrozo, 2020, for some key contributions in this direction).

low chance. Hawthorne and Lasonen-Aarnio (2009) note that a more sophisticated version of this constraint is the first step on a path towards future skepticism (which, like me, they view as undesirable).[2] These challenges do not require commitment to the idea that the future is any different from the past in either ontology or structure. That is to say: they arise even for those who believe that that future contingents are fully bivalent, and that there is no indeterminacy with regards to which future is ours. All that is needed is the belief that incompatible propositions about the future may have non-extreme chances.

Though this work at the interface between the theory of chance and the theory of knowledge is of prime importance, my focus will be on a nearby issue. I concentrate on challenges that arise when we adopt—even just for the sake of intellectual exploration—the idea that the future is open. Very roughly, the open future hypothesis is the claim that future events and states of affair are unsettled, or indeterminate.[3] This indeterminacy is emphatically not epistemic. It also goes beyond the claim that multiple ways the future might unfold have positive chance. The next section expands upon the content of this hypothesis.

My central question is whether proponents of the open future hypothesis must be future skeptics. At first sight, the answer would appear to have to be "yes", for reasons broadly related to the factivity of knowledge — the widely accepted idea that one can only bear the knowledge relation to true propositions. If knowledge is factive, knowing that the crystal glass will shatter entails the proposition that the crystal glass will shatter. If this proposition is indeterminate, it would seem to follow that one cannot know it. I will eventually question this last step. However, it might seem hard to see how to avoid that step if one thought it a consequence of the open future that the principle of bivalence for future contingents fails. Fortunately, as I clarify in section 2, we can develop an account of foreknowledge against an attractive metaphysical background which is compatible with both belief in the open future and endorsement of bivalence.

Once this metaphysical backround is in place, my proposal is that, to stop the slide from the open future hypothesis to future skepticism, we must entertain the idea that it might also be indeterminate whether one is in a knowledge state.

[2] For more on this debate, in addition to Hawthorne and Lasonen-Aarnio (2009), see the reply by Williamson (Williamson, 2009). Related puzzles at the interface between the theory of knowledge and the theory of objective chance are discussed by Dorr *et al.* (2014), Goodman and Salow (2018, ms.), Stalnaker (2019, esp, pp. 146-148, 238).

[3] For a variety of (incompatible) contributions on the open future, see Łukasiewicz (1970); Thomason (1970); Belnap *et al.* (2001); MacFarlane (2003, 2014); Hirsch (2006); Barnes and Cameron (2009, 2011); Torre (2011); Correia and Iacona (2013); Cariani and Santorio (2018); Torrengo and Iaquinto (2020).

Specifically, if it is indeterminate that some future contingent proposition is true, then it can be indeterminate whether an agent knows that proposition. For example, if it is indeterminate whether the crystal glass will break, it can be indeterminate whether I know that it will break.

This paper aims to clarify this view, by investigating the formal contours of a concept of indeterminate knowledge that is suitable to play this role. After presenting some background (sections 2 and 3), I consider three ways of arguing for future skepticism (section 4). These three arguments depend on three constraints on the relation between indeterminacy and knowledge, each of which leads to future skepticism by means of straightforward (classical) logic. The need to avoid these constraints nudges us towards entertaining indeterminate knowledge states (section 5 and 6). Sections 7 and 8 expand on some specific issues that arise once we do entertain them. The final section outlines a standard model theoretic analysis of the formal principles on the interaction between indeterminacy and knowledge that are discussed in the course of the paper.

Although I will be theorizing under some version of the open future hypothesis, I am myself neither committed to it, nor to its rejection. I take it as a central point of methodology that we can make progress by assuming each relevant hypothesis provisionally, exploring the plausibility of its epistemic implications, and revising our confidence in each hypothesis accordingly.

2 Background on the open future hypothesis

It is notoriously controversial how best to characterize the open future hypothesis (Torre, 2011). The idea of the open future has historically been associated with a cluster of metaphysical and semantic theses we should not necessarily take aboard.

Branching. The indeterminacy of the future is well represented by a conception on which possible worlds overlap and branch (Belnap and Green, 1994; Belnap *et al.*, 2001).

Asymmetric indeterminism. The past is nomically necessary but the future is not.

Non-bivalence. Future contingent propositions are neither true nor false. (Łukasiewicz, 1970; Thomason, 1970; MacFarlane, 2003)

The first two theses are not directly relevant here (except insofar as they can be used to argue for the third) and I will set them aside. For the record, I think both are problematic if understood as characterizations of the openness of the future.[4]

If we are to avoid skepticism, it is absolutely critical however that we reject the last assumption, for it is one of the main drivers of the slide from belief in the open future to skepticism. If knowledge requires truth and future contingents are neither true nor false, then future contingents are unknowable.

I will build my resistance strategy against the background of a cluster of metaphysical views associated with Barnes and Cameron (2009, 2011) and with a broader stance in support of bivalent indeterminacy.[5] For Barnes and Cameron, the indeterminacy of the future amounts to the thesis that there are multiple ways the future might unfold are metaphysically on a par at any given moment — that nothing settles which of multiple candidates is the actual future. Though this picture allows for indeterminacy, it divorces the question of indeterminacy from the question of bivalence. It is determinate that the proposition that the glass will shatter is either true or false. At the same time, however, that proposition is not determinately true and not determinately false.

An analogy with certain theories vagueness can help conceptualize this situation: suppose I am located on the border between Washington DC and Maryland. It is coherent to describe the situation by saying that I am determinately either in DC or in Maryland, but not determinately in DC and not determinately in Maryland. In particular, it seems plausible that I am not located in some third location that's neither DC nor Maryland. Adopting a similar stance with respect to modeling the semantic consequences of the indeterminacy of the future opens up a path towards a non-skeptical stance. Having registered this analogy, I will also note that my development depends on assumptions that are distinctive of the case of the future and do not generalize to problems of vagueness.

3 A formal language for indeterminacy and foreknowledge

Let us set up a formal language in which to talk about knowledge and determinacy. I adopt the convention that roman sans-serif letters (A, B, C) are metalinguistic

[4] For critiques, see Barnes and Cameron (2011), Cameron (2015, ch.5), Cariani (2021, section 11.2).

[5] In addition to Barnes and Cameron, this position is entertained in Hirsch (2006) Cariani and Santorio (2018), Cariani (2021, chs.10-11). These are not the only attempts to reconcile bivalence and the openness of the future. There are related approaches within the broad family of "Thin Red Line" views (Malpass and Wawer, 2012; Borghini and Torrengo, 2013; Iacona, 2013; Torrengo and Iaquinto, 2020). These views are not equivalent to the bivalent indeterminist view I adopt in the paper, but they might also be leveraged to support the present epistemological development.

variables ranging over sentences of this language. By contrast, bold letters (e.g. $\mathbf{A}, \mathbf{B}, \mathbf{C}$) range over propositions, with the further specification that, in context, I use $\mathbf{A}$ to refer to the proposition expressed by sentence A. Quasi-quotation is implied whenever these metavariables are involved in complex strings, but never represented.

The primitive grammar of the language is:

$$\mathsf{A} ::= p \mid \neg\mathsf{A} \mid (\mathsf{A} \wedge \mathsf{A}) \mid K\mathsf{A} \mid D\mathsf{A},$$

This is a modal language. In particular, that the K and D operators are box-like modal operators and can embed freely within the language.[6] Disjunction ($\vee$) and the material conditional ($\rightarrow$) are defined as usual — respectively: $\neg(\neg\mathsf{A} \wedge \neg\mathsf{B})$ and $\neg(\mathsf{A} \wedge \neg\mathsf{B})$. We assign to D and K corresponding dual operators — $\lozenge$ for D and $\blacklozenge$ for K. Thus, $\lozenge$ captures what is possible in light of the determinate facts, while $\blacklozenge$ captures what is possible in light of the knowledge encoded in K.

There is an important difference with regard to the temporal anchoring of these operators. K and $\blacklozenge$ model the knowledge state of specific agent at a specific moment in time. I do not make the same assumption about D and $\lozenge$, because we need these to track determinacy facts dynamically (details to come in sections 5 and 9). To reflect this difference with terminological distinction, I speak of a *designated* moment for K and of a *relevant* moment for D. Let us illustrate this with a mixed sentence that will play an important role in the following: $\lozenge K\mathsf{A}$. To make this more concrete, suppose that the value of A is the sentence *it will rain tomorrow* and the designated moment for K is Monday, while the relevant moment for D is Tuesday. According to the interpretation strategy I just outlined, $\lozenge K\mathsf{A}$ means that the determinate facts on Tuesday are compatible with one's having known $\mathbf{A}$ on Monday. Also according to the same interpretation strategy, the relevant time for D might change, but the designated time for K does not.

In talking about the possibilities that are tracked by these operators, it is convenient to bend some standard terminology. Say that D and $\lozenge$ quantify over *historical* possibilities, while K and $\blacklozenge$ quantify over *epistemic* possibilities. To say that it is historically possible that my daughter will lose her front tooth tomorrow is to say that the determinate facts at the present time do not settle that she will not lose her front tooth tomorrow. Associating determinacy operators with the notion of historical possibility signals that the only source of indeterminacy to

[6] As Barnes and Williams (2011, §6) highlight, allowing determinacy operators to embed freely under other operators makes our analytical task significantly more complicated than if we only allowed D to occur as the main operator of formulas.

be considered is the indeterminacy of the future. This usage is not congenial to those who maintain that the indeterminacy of the future belongs to a broader conceptual category that includes other kinds of indeterminacy (as suggested in Barnes and Williams 2011 and also elaborated by Williams 2012, 2014). Although I do think that the indeterminacy of the future, I reiterate here that I make these terminological choices for convenience.

Likewise, my talk of epistemic possibilities differs from the prevalent usage of the phrase. In particular, the prevalent usage of the phrase "epistemic possibilities" refers to possibilities that are tracked by natural language epistemic modals such as the English *might*. My use of "epistemic possibility" here is technical and it picks out whatever possibilities we must quantify over in our (Hintikka-style) model of knowledge. In other words, despite the terminology, the symbol $\Diamond$ and the English *might* have no interesting semantic connection.

Having started out with a language with two modal operators K, and D, we can use the latter to define an indeterminacy operator I:

$$I\mathsf{A} := \neg D\mathsf{A} \ \& \ \neg D\neg\mathsf{A}$$

A proposition is indeterminate in truth value if and only if neither it nor its negation are determinate. Given the duality assumption, $I\mathsf{A}$ is also equivalent to $\neg D\mathsf{A} \ \& \ \Diamond\mathsf{A}$ (i.e., A is not determinately true but it is historically possible), and perhaps most perspicuously to $\Diamond\neg\mathsf{A} \ \& \ \Diamond\mathsf{A}$ (both A and its negation are historically possible).

4 Formal correlates of future skepticism

The future skeptic maintains that indeterminacy is a barrier to knowledge. Some defenders of the open future hypothesis embrace the skeptical connection between indeterminacy and unknowability. Even Barnes and Cameron were initially tempted by this position:

> [T]he unknowability [of future contingents] on our account is not constitutive of the indeterminacy, but rather a consequence of it. Because the truth-value of the propositions expressed by presently uttered future-directed sentences are metaphysically unsettled, such propositions cannot be known—such propositions have a truth-value, but their truth-value must remain epistemically inaccessible until the unfolding of the future settles which truth-value they in fact have (Barnes and Cameron, 2009, p. 298).

Barnes and Cameron revised this stance in §4.4 of Barnes and Cameron (2011), where they briefly sketch a view that is very much like the one I develop here.

The present focus, however, is on articulating the skeptical position with more precision and depth. We want to represent in the formal language the thesis that no one knows indeterminate propositions. Here is its natural implementation:

Skeptical core. $I\mathsf{A} \to \neg K\mathsf{A}$

Recall that $I\mathsf{A}$ unpacks as $\Diamond\mathsf{A}\ \&\ \neg D\mathsf{A}$. Of these two conjuncts, only the second, $\neg D\mathsf{A}$, is relevant to support the skeptical consequent, $\neg K\mathsf{A}$. In other words, what matters to our inability to know **A** is the fact that it is not determinately true.

Though importantly different, the skeptical core fits with the spirit of Field's (2000) "rejectionist" analysis of belief in indeterminate contents. Writing with vagueness as the central case study, Field argues that accepting that a proposition is indeterminate involves absolute rejection of both it and its negation. These positions share the idea that indeterminacy is a barrier to certain cognitive standings, be they knowledge or positive credence. Of course, there are important, and ultimately irreconciliable differences between these approaches to indeterminacy.

It is important to note some standard ways of deriving the skeptical core from basic classical logic, together with some additional assumptions. Our first purported proof is a warm-up:

1. $\vdash K\mathsf{A} \to \mathsf{A}$ — factivity of K
2. $\vdash \mathsf{A} \to \neg I\mathsf{A}$ — truth-to-determinacy
3. $\vdash K\mathsf{A} \to \neg I\mathsf{A}$ — 1,3 transitivity of $\to$
4. $\vdash I\mathsf{A} \to \neg K\mathsf{A}$ — 4, contraposition for $\to$

Open futurists come prepared to this challenge: they must reject the idea that truth entails determinacy. No indeterminist needs to accept that if the coin lands heads, then it is determinate that it lands heads.

There is an important choice point concerning how to deny step 2. A non-classical approach, leaning on the concept of global validity in supervaluational semantics (see e.g. see Williamson 1994 ch. 5, and Varzi 2007), has it that the argument with premise A and conclusion $\neg I\mathsf{A}$ is valid (i.e. $\mathsf{A} \vdash \neg I\mathsf{A}$), and yet that it does not follow that $\mathsf{A} \to \neg I\mathsf{A}$ is a theorem. In other words, the deduction theorem fails. The classical alternative would have both the theorem and the premise/conclusion argument fail. The model theory I present in section 9 implements this classical blueprint. However, the non-classical alternative is a well established and in no way ruled out by anything I say here.

A second purported skeptical proof goes via a different kind of insight. One might suppose that if it is not determinate which of two worlds is mine, then my epistemic state cannot distinguish between them. If w and v have exactly the same history but divergent futures then my epistemic state lacks the discriminatory capacity to identify one of these worlds as mine. More explicitly, if a possibility is historically open, I am not in a position to rule it out.

1. $\vdash \Diamond\neg \mathsf{A} \rightarrow \blacklozenge\neg \mathsf{A}$ — skeptical bridge
2. $\vdash I\mathsf{A} \rightarrow \blacklozenge\neg \mathsf{A}$ — 1, antecedent strengthening for $\rightarrow$
3. $\vdash \blacklozenge\neg \mathsf{A} \rightarrow \neg K\mathsf{A}$ — duality of $\blacklozenge$ and K
4. $\vdash I\mathsf{A} \rightarrow \neg K\mathsf{A}$ — 2,3, transitivity of $\rightarrow$

It is sometimes helpful to restate the skeptical bridge purely in terms of the dual modals, i.e. as $\neg D\mathsf{A} \rightarrow \neg K\mathsf{A}$.

The skeptical bridge bears some resemblance to principles that are the focus of the literature on knowledge and objective chance. Following Hawthorne and Lasonen-Aarnio (2009), the relevant connection is typically forged through some version of a safety requirement on knowledge. Roughly speaking, the safety requirement says that one's knowing **A** in w requires that **A** be true at all the closest worlds in which one believes it. In this kind of framework, a close analogue to the skeptical bridge might say that historical possibilities are always close in the sense that is relevant for safety.

A third type of proof relies on the principle that knowledge is always determinate. This is explicitly posited by Belnap *et al.* (2001) who say "when we have knowledge, it is settled true that we have it" (p. 56). Since their concept of settled truth is close to the present concept of determinacy, this thought can be approximately represented by the principle that $K\mathsf{A} \rightarrow DK\mathsf{A}$. This principle also entails the skeptical core.

1. $\vdash K\mathsf{A} \rightarrow DK\mathsf{A}$ — determinacy of knowledge
2. $\vdash K\mathsf{A} \rightarrow \mathsf{A}$ — axiom T for K
3. $\vdash D(K\mathsf{A} \rightarrow \mathsf{A})$ — 2, necessitation for D (in **LKD**)
4. $\vdash DK\mathsf{A} \rightarrow D\mathsf{A}$ — 3, axiom K for D + modus ponens rule

5. $\vdash K\mathsf{A} \to D\mathsf{A}$ 1,4, transitivity of $\to$

6. $\vdash \neg D\mathsf{A} \to \neg K\mathsf{A}$ 5, contraposition for $\to$

7. $\vdash I\mathsf{A} \to \neg K\mathsf{A}$ 6, antecedent strengthening for $\to$

These proofs might not convince an anti-skeptic, but they demarcate boundaries within which an anti-skeptic open-futurist must operate.

The skeptical bridge and the determinacy of knowledge are tightly related. The determinacy of knowledge entails the skeptical bridge.[7] In section 9, with a model theory in hand, we will be able to see that the entailment is not symmetric absent any other assumptions. However, it can be strengthened to equivalence by assuming other principles. For example, the KK principle would suffice for this task.[8]

Before I get to what that positive view must look like, it is important to take note of a slight strengthening of the skeptical core. This is the principle:

Strengthened core. $I\mathsf{A} \to \neg\Diamond K\mathsf{A}$

The difference between the skeptical core and its strengthening is somewhat subtle. If we think in terms of possible worlds, we can put the difference as follows. The core says that **A**'s indeterminacy in w is incompatible with **A** being known in w. The strengthening says that **A**'s indeterminacy in w is incompatible **A**'s being known in w and also in any world other world that is exactly like w as far as the determinate facts go.

Within a standard relational possible world semantics (see again section 9), the core does not entail its strengthening. Figure 1 shows a countermodel: $\neg Dp$ is true at w and so is $\neg Kp$. However, at that world it is historically possible that one is in v, where p is determinate and known. These countermodels can be ruled out by stipulating that the relation underlying D is Euclidean (i.e., whenever wR_Dv & wR_Dz, we must have vR_Dz). Thus the core and its strengthening are equivalent modulo the axiom that characterizes the class of Euclidean frames ($\neg D\mathsf{A} \to D\neg D\mathsf{A}$).[9]

[7] Here is a proof sketch. If K is factive, $\vdash K\mathsf{A} \to \mathsf{A}$. Since D is a normal modal operator, satisfying the necessitation rule and axiom K, $\vdash DK\mathsf{A} \to D\mathsf{A}$. Chaining this with $\vdash KA \to DKA$, we get $\vdash K\mathsf{A} \to D\mathsf{A}$. By contraposition $\vdash \neg D\mathsf{A} \to \neg K\mathsf{A}$ which after appropriate substitutions based on duality principles yields $\vdash \Diamond\neg\mathsf{A} \to \blacklozenge\neg\mathsf{A}$.

[8] Contraposing the skeptical bridge and using duality, we get $K\mathsf{A} \to D\mathsf{A}$. Consider the class of instances of this principle in which A is replaced with $K\mathsf{A}$ to get $KK\mathsf{A} \to DK\mathsf{A}$. If KK holds, we have $K\mathsf{A} \to KK\mathsf{A}$, which can be chained with that last principle to yield, $K\mathsf{A} \to DK\mathsf{A}$.

[9] Start with $\neg D\mathsf{A} \to \neg K\mathsf{A}$. If this is a theorem, so is its necessitated version $D(\neg D\mathsf{A} \to \neg K\mathsf{A})$. Distributing D across the conditional yields $D\neg D\mathsf{A} \to D\neg K\mathsf{A}$, i.e., $D\neg D\mathsf{A} \to \neg\Diamond K\mathsf{A}$. Chaining axiom 5 with this theorem yields the strengthened core. If one accepts the standard relational model theory for normal modal logics, it is easy to see that without it, the strengthened core fails.

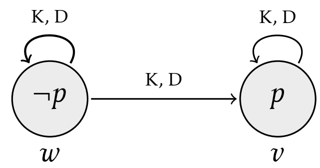

Figure 1: The gap

Another way to derive the strengthened core is by accepting appropriate strengthenings of the principles deployed in earlier proofs, in particular:

Strengthened bridge. $\neg D\mathsf{A} \to \neg\Diamond K\mathsf{A}$

Strengthened determinacy of knowledge. $\Diamond K\mathsf{A} \to DK\mathsf{A}$

Although there is logical daylight between the core and its strengthening, there is little reason for the skeptic to go for the core but reject the strengthening. For one thing, many conceive of the indeterminacy of the future so as to support the Euclidean property for R_D. More importantly, however, the view that accepts the core but denies the strengthening is only viable as a strange mix of skeptical and anti-skeptical features. It deviates from full-on skepticism by imagining that is never known in any world in which it is not determinately true. However, it is known in some worlds that count as historical possibility from the perspective of the base world. This mix of skepticism and anti-skepticism is not any kind of middle ground. It is just satisfactory no one. For this reason, I take the stable skeptical position to be committed to the stronger principles.

5 The anti-skeptical baseline

The need to reject the determinacy of knowledge points to an approach that allows us to avoid skepticism while retaining a broadly open-futurist framework. Short of denying some of the classical assumptions involved in the third proof, the *only* path for an anti-skeptical open futurist is is to entertain the idea that it can sometimes be indeterminate whether one is in a knowledge state with respect to an indeterminate proposition.

In Cariani (2021, ch. 10-11), I consider an analogous approach to related problems about assertion. Suppose, for the sake of argument, that assertion is subject to a truth norm—one may assert **A** only if **A** is true. The assertion analogue of skepticism is the view that if it is indeterminate a proposition is true, then that proposition is determinately not assertible (MacFarlane, 2014, ch. 9). The problem with this position is that intuitively we are in a position to make a great many perfectly legitimate assertions about the future.[10] Instead of rejecting the whole class of future contingents as unassertible, it is more plausible to adopt a more permissive position, according to which it is itself indeterminate whether the truth norm is met by assertions of indeterminate propositions. The structural point is that indeterminacy in whether some factual requirement is met lifts up to indeterminacy of normative statuses that depend on that factual requirement being met. For parallel reasons, I claim, indeterminacy of truth may sometimes result in indeterminacy of knowledge.

But in what sense does appeal to indeterminate knowledge states get us out of skeptical danger? Addressing this concern requires identifying some relevant ways in which indeterminate knowledge states are different from states of determinate ignorance.

The most distinctive aspect of the indeterminacy posited by open future theorists is that the cloud of indeterminacy shrinks as we progress through time. Here are Barnes and Cameron making just this point:

> [W]e hold that whilst future contingents have no determinate truth-value, they are, determinately, either true or false. Furthermore, we hold that the truth-value they have can be revealed as time progresses and events unfold. So suppose on Friday you make a prediction: 'Aliens will invade tomorrow'. Your utterance lacks a determinate truth-value. But come Saturday, when aliens are mercifully absent, you can look back and say that your prediction was false. (Barnes and Cameron, 2011, p. 4)

This dynamic nature seems uniquely specific to the indeterminacy of the future. It is not shared by other varieties of indeterminacy that are invoked to understand vagueness and most varieties of paradox.

Let us illustrate it with a specific example. Suppose that on Monday, the historically possible worlds O_M are agnostic about the weather on Tuesday. Then, on Tuesday, it snows. At that point, the set of historically possible worlds gets trimmed down to $O_T \subset O_M$ consisting of all the worlds that share the history up to Monday, but additionally feature snow on Tuesday. Let A be some sentence stating in our object language the claim that it snows on Tuesday. Then A might go from

[10] See Besson and Hattiangadi (2014) and (Cariani, 2021, ch.10). In addition, Cariani (2020) and (Cariani, 2021, ch.9) argue against the idea that there are no future-directed assertions. According to these are views, we perform speech acts of prediction that are subject to different standards than ordinary assertion. Against this, I argue that most predictions in fact are assertions.

indeterminate on Monday, to determinate on Tuesday and thereafter (as shown in Figure 1). It is this dynamic element distinguishes the indeterminacy of the open future from other kinds of indeterminacy.

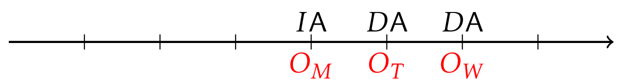

Figure 2: Dynamic indeterminacy

To illustrate how knowledge interacts with this indeterminacy let us make a simplifying assumption. Assume that there is a state a thinker can be in, which is just like the state of knowing **A** minus the truth of **A**. Call this state *pre-knowledge* of **A**. We may well doubt that such a state exists, for reasons roughly related to Williamson's (2000, ch. 2, 3) critique of attempts to factorize knowledge in internal and external components. However, I do not suggest taking this concept all that seriously: we just help ourselves to it for illustration's sake.[11] Moreover, there are some nearby states whose existence we have no reason to doubt, such as historically possible knowledge (what would be expressed by $\Diamond K$) and epistemically possible knowledge (corresponding to $\blacklozenge K$).[12]

Suppose that on Monday one pre-knows **A** but **A** is indeterminate ($O_M \not\subseteq \mathbf{A}$ and $O_M \not\subseteq \overline{\mathbf{A}}$); then it is indeterminate on Monday whether one knows **A**. On Tuesday **A** is no longer indeterminate ($O_T \subseteq \mathbf{A}$). Here is the central thought: from Tuesday's perspective, it is determinate that I did know (on Monday) that it would snow (on Tuesday). It is important that this analysis need not involve commitment to semantic relativism about future contingents (or about determinacy claims) in the style of MacFarlane (2003, 2014), though it is compatible with it.

This dynamic component is absent from states of determinate ignorance. If on Monday it is determinately the case that I am ignorant about Tuesday's weather, there is no possible future development in which my Monday state can correctly be described as knowledge. What is possible, of course, is that *on Tuesday* I occupy

[11] In fact, it is not clear whether Williamson's arguments rule out the possibility of pre-knowledge — that is of a state that is just like knowledge except for the truth of the matter, especially against a background that allows for indeterminate truth. Strictly speaking Williamson argues that there is no sharp way of separating internal conditions involved in knowledge from external ones. In other words, that there is no pure inner "factor" of knowledge. But this is compatible with knowledge involving some mix of internal conditions, some purely external conditions (such as truth), and some 'prime' conditions.

[12] For discussion of these principles, and proposals to think of epistemically possible knowledge as providing account of justification, see Stalnaker (2006, 2019); Bird (2007); Ichikawa (2014); Carter and Goldstein (2020).

a new state which counts as knowledge that it is snowing on Tuesday. But that is hardly news and hardly of relevance to a debate with the future skeptic.

An important question is how much of the work that we need human foreknowledge to do can be done by appealing to states of indeterminate knowledge. In my view, dynamically evolving states of indeterminate knowledge would allow us to avoid many pernicious consequences of future skepticism.

Let me illustrate by considering the role that knowledge plays in discourse. According to a widespread (but by no means universally accepted) view, assertions are governed by a knowledge norm: there is something defective about asserting **A** when one does not know **A**.[13] Suppose that on Monday I asserted the proposition that it will snow on Tuesday. Suppose again that my state on Monday is pre-knowledge of that proposition and that the proposition is indeterminate in truth-value. Then it is indeterminate whether I know what I asserted, and thus indeterminate whether I have met the norm of assertion. That indeterminacy however is not sufficient to warrant the claim that my assertion is defective in the ways that assertions that determinately violate a norm are. Indeed on Tuesday, as the indeterminacy is resolved, it is determinate that I knew that it would snow, and thus that my Monday assertion did meet the knowledge norm of assertion.

6 Lifting principles

It would be nice to have a principle that captures the connection between indeterminacy and knowledge. For obvious reasons, that principle cannot be:

Unrestricted lift. $I\mathsf{A} \rightarrow IK(\mathsf{A})$

Indeterminacy only lifts in this way when one is in the right mental state towards a proposition. In our colorful but non-serious language, indeterminacy in the truth-value of **A** only lifts when one pre-knows **A**. After all, if indeterminacy alone were sufficient for lifting, there would be no contingent propositions about the future of which any one of us is determinately ignorant, and instead there are plenty.

To patch this problem with the unrestricted lift without explicitly appealing to pre-knowledge states in our official theory, we stipulate:

[13] Williamson (2000) is often associated with this idea, but Williamson adds in the idea that this norm is *constitutive* of the speech act of assertion. A much larger swath of theorists would accept the weaker claim that knowledge is a normative requirement of non-defective assertions. Indeed, this is plausibly a consequence of Grice's (1975) Maxim of Quality.

Restricted lift. $\Diamond K\mathsf{A} \to (I\mathsf{A} \to IK\mathsf{A})$

The proposal is that indeterminacy lifts whenever the settled facts up to the designated time leave open whether I know **A**.

Under basic assumptions, the restricted lift follows from the standard way of combining the logics of knowledge and indeterminacy. In particular, suppose that the isolated logics of knowledge and determinacy are normal modal logics. More precisely, suppose that $\mathcal{L}_K \subseteq \mathcal{L}$ is the modal language with K as its only operator and $\mathcal{L}_D \subseteq \mathcal{L}$ is the modal language with D as its only operator. Next, say that the isolated logic of K (call it '**LK**') and the isolated logic of D (call it '**LD**') are both normal modal logics in their respective languages. The *fusion* of two normal modal logics is ordinarily defined as the smallest normal modal logic that extends both (Kurucz, 2007; Carnielli and Coniglio, 2020).

There are some important facts to be highlighted about the fusion **LKD** of **LK** and **LD**. These facts do not depend on the strength of the isolated logics. (To substantiate this point, I will generally work with whichever logic makes the task at hand harder.) A first observation is that the skeptical principles from section 4 are not theorems of **LKD**. That is, the skeptical core, $I\mathsf{A} \to \neg K\mathsf{A}$, the skeptical bridge, $\Diamond\neg\mathsf{A} \to \blacklozenge\neg\mathsf{A}$, and the determinacy of knowledge, $K\mathsf{A} \to DK\mathsf{A}$, are not theorems of **LKD**, even if the isolated logics **LK** and **LD** are each taken to be as strong as S5. For now, I state these claims without proof, but a proof is implied in the model theoretic characterizations of section 9. By contrast, the restricted lift, $\Diamond K\mathsf{A} \to (I\mathsf{A} \to IK\mathsf{A})$, is a theorem of **LKD**, even if the isolated logics are each as weak as K.[14]

It is somewhat surprising that the restricted lift is valid in **LKD**, because logical fusion is largely neutral on the choice between skeptical and non-skeptical approaches. The key to understand this fact is to note that the restricted lift is endorsed by the skeptic as merely *trivially* valid. Recall the strengthened skeptical

[14]To see why the restricted lift is a theorem of **LKD**, start by recalling that $I\mathsf{A}$ unpacks to $\Diamond A \;\&\; \neg D\mathsf{A}$. In particular, the consequent of the restricted lift, $IK\mathsf{A}$, unpacks to $\Diamond K\mathsf{A} \;\&\; \neg DKA$. The first conjunct of $\Diamond K\mathsf{A} \;\&\; \neg DKA$ can be proven directly from the assumption $\Diamond K\mathsf{A}$ — which formula is one of the antecedents in the restricted lift. In light of that, the restricted lift is established as soon as we have:

(i) $I\mathsf{A} \to \neg DK\mathsf{A}$

In turn, (i) follows from a basic theorem of the fusion logic **LKD**:

(ii) $\neg D\mathsf{A} \to \neg DK\mathsf{A}$

This is a theorem in the combined logic since $K\mathsf{A} \to \mathsf{A}$ is a theorem of **LK** and thus a theorem of **LKD**. By necessitation, $\vdash D(K\mathsf{A} \to \mathsf{A})$ which after trivial manipulations entails (ii).

core: $I\mathsf{A} \rightarrow \neg\Diamond K\mathsf{A}$. Given this, the antecedent of the outer conditional of the restricted lift, $\Diamond K\mathsf{A}$, is incompatible with the antecedent of the inner conditional, $I\mathsf{A}$. The local upshot is that the skeptic's endorsement of the validity of the restricted lift has nothing to do with endorsing indeterminate knowledge as a genuine possibility.

The broader upshot is that, although the restricted lift is a natural expression of a view that the antiskeptic endorses, it would be incorrect to suppose that it *characterizes* the antiskeptical position. In fact, I conjecture that in this language there is no way of characterizing the anti-skeptical position in terms of the validity of any one statement.[15] As a result of these considerations, it is best to view the dispute between the skeptic and their opponent as entirely centering on the former's acceptance of the skeptical core (and its cognates) and the latter's rejection of these validities, coupled with some claim to the effect that the restricted lift is taken by the anti-skeptic to be *non-trivially* valid.

7 Indeterminacy and the unidirectionality of knowledge

We saw that future skeptics are committed to principles that exceed the fusion of the isolated logics of determinacy and knowledge. As it turns out, those who want to avoid future skepticism also face important choice points that require enriching the fusion logic. In this section, I take up one such example, first by presenting an informal idea and then by matching it with a formal constraint.

The question to be considered involves constraints on the distribution of indeterminate knowledge states across partitions of future possibilities. Suppose I buy Linne a gift. When she gets it, she will be *happy*, *sad*, *indifferent*, or *undecided*. These are mutually exclusive and exhaustive of the future possibilities. Notate the partition corresponding to Linne's responses as $\{\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3, \mathbf{A}_4\}$. In such a situation, we may wonder whether it is possible for one to have indeterminate knowledge of two or more of the $\mathbf{A}_i$'s. That is, we'd like to know if there are x and y such that:

$$IK\mathsf{A}_x \ \& \ IK\mathsf{A}_y$$

There is good reason to rule this out and adopt instead the view that the kind of indeterminate knowledge of interest here is by and large *unidirectional*. That is to say, with a small class of possible exceptions to be considered shortly, an ordinary

[15] We could get around this problem if our language had propositional quantification. In that setting we could directly express the claim that there are propositions that are indeterminately known. I do not pursue the complications involved with this move here.

human agent cannot indeterminately know a proposition and also indeterminately know another, incompatible proposition.

One motivating intuition for unidirectionality might involve the idea that indeterminately knowing **A** is a state that is as similar as possible to knowing **A**, except for the unsettled truth status of the proposition. And just as one cannot know incompatible propositions, one cannot indeterminately know incompatible propositions.

One way of building on these points is to reflect on consistency constraints on potential requirements of knowledge — e.g. on justification. Suppose one believes that Linne will be sad and also believes that she will be indifferent (recall that we are assuming these states are incompatible). There is some plausibility to the thought that those contradictory beliefs generate a conflict that would strike down at least one, if not both, of the beliefs from counting as justified, and thus as, if justification is required for knowledge. The fact that they are logically incompatible suggests that, for each belief, one accepts a proposition that defeats it. If both defeaters have equal strength, then both knowledge items have been defeated. In the alternate case in which one of the defeaters should be prioritized over the other, it still determinate that one does not know both $\mathbf{A}_x$ and $\mathbf{A}_y$.

Not everyone will find this argument persuasive. To start, there is conceptual daylight between accepting a proposition that could count as defeater for a belief, and having the belief be defeated. Perhaps one's belief is defeated only if one is aware of the defeat relation. For a related theoretical point, considerations related to the preface paradox (as well as other paradoxical scenarios) convince some theorists that there can be simultaneous justification for beliefs in incompatible contents. This is not the place to explore, much less resolve, this dialectic.[16] I just draw the moral that some might find unidirectionality attractive because of consistency requirements on justification, while others might not. Should this defense of unidirectionality were to fail, there are other paths to supporting unidirectionality of indeterminate knowledge. In particular, we might sign up for a safety-based or normality-based account of knowledge (in the style of Smith 2016; Goodman and Salow 2018, ms.; Carter and Goldstein 2020). It seems plausible to ground the unidirectionality of knowledge in the unidirectionality of certain important normality relations.

[16] Foley (1979) denies consistency constraints on justification. For other defenses of the idea that there can be justified or rational inconsistency, see Christensen (2004); Worsnip (2016); Littlejohn and Dutant (2019); Staffel (2019). On the other side, Neta (2018) responds to Worsnip's argument. For examples of accounts that support such consistency constraints, see Pollock and Cruz (1999); Smith (2016). Also, the analyses of justification in terms of knowledge mentioned in footnote 12 generally support consistency constraints.

There is another limit to my argument for the unidirectionality of indeterminate knowledge. The argument — and in fact the principle of unidirectionality itself — might fail, unless we assume in the background that the only relevant source of indeterminacy is the truth status of the relevant propositions. Some philosophers working on vagueness and paradox (Dorr, 2003; Barnett, 2009; Caie, 2012) allow that one might sometimes be in a state of indeterminate belief. Suppose one can be in a state of indeterminate belief towards the proposition that Linne will be happy ($\mathbf{A}_1$) and in a state of indeterminate belief towards the proposition that she will be sad ($\mathbf{A}_2$). This might hold while it is determinate that I have exactly one relevant belief. It seems coherent to enrich this picture by imagining that there is one resolution of the indeterminacy on which $\mathbf{A}_1$ is true, believed and known, and a second resolution of the indeterminacy on which $\mathbf{A}_2$ is true, belied and known. Such a situation might allow multi-directional indeterminate knowledge *via* multi-directional indeterminate belief. That is, it would allow that it is determinate that I know one of the $\mathbf{A}_i$'s but indeterminate which.

This suggests that unidirectionality might need to be restricted. After all, if one entertains the possibility of indeterminate beliefs, one entertains the possibility of indeterminate beliefs about the future. This is no refutation of the principle of unidirectionality, since we can restrict it to determinate beliefs. So revised, it states that, as long as we assume that all the relevant states of beliefs are determinate, then states of indeterminate knowledge ought to be unidirectional. This restriction would be problematic if it trimmed out something of present interest, and in particular if there was a reason to assume that the indeterminacy *of the future* should by itself (that is: independently of vagueness or paradox) lead us to entertain indeterminate belief states. But there is no reason to assume that. Our state of belief in non-vague, non-paradoxical future contents do not seem to be relevantly different from similar states of belief in non-vague, non-paradoxical contents that are about the past and the present. The intuitive pull one feels towards appealing to indeterminate belief for vague contents is completely absent when it comes to the indeterminacy of the future.[17]

[17] This discussion allows us to go on a brief but interesting detour on the philosophical significance of unidirectionality. The unidirectionality principle is important for distinguishing human foreknowledge from the kind of foreknowledge one might consider assigning to an omniscient being (e.g. to the God of the Abrahamic religions). More specifically, Todd and Rabern (2021) explore (and ultimately reject) the idea that we might make sense of divine omniscience in terms of a concept of indeterminate knowledge. However, in the context of the debate on divine foreknowledge, it is very important that indeterminate knowledge be *omnidirectional*. That is, if Π is a partition of the historical possibilities, one would need it to be the case that for each $\mathbf{A} \in \Pi$, God indeterminately knows $\mathbf{A}$. The reason to appeal to indeterminate knowledge stems from a desire to square the idea of an open future with the idea that God is omniscient. After all, if it is objectively indeterminate

The natural implementation of unidirectionality in our formal setup is in terms of the following principle.

Unidirectionality. $(IK\mathsf{A}\ \&\ IK\mathsf{B}) \rightarrow \Diamond(\mathsf{A}\ \&\ \mathsf{B})$

Informally: one can have indeterminate knowledge towards two propositions **A** and **B** only when it is historically possible for both **A** and **B** to hold. If we wanted to restrict this principle, so that it only holds under the assumption of determinate belief, we could consider expanding the formal language by adding belief operators and restricting the principle to those cases in which one determinately believes the two propositions. Unidirectionality does not follow from the mere fusion of **LK** and **LD** (This will be proven in section 9). In fact, unidirectionality is a paradigm example of a substantive principle that needs to be added on to flesh out the anti-skeptical open futurist's world view.[18]

8 Indeterminate knowledge and inference

Once we grant that our state towards some future contingent propositions is indeterminate knowledge, a question arises as to whether it is possible to be in this state towards a determinate truth. Suppose in particular that one infers a determinate truth **C** from a premise **A** which is not determinately true and not determinately false. Suppose that the state one is in towards **A** qualifies as indeterminate knowledge. What should we say about the status of one's belief towards **C** in such a situation?

It is tempting to say that beliefs that are based entirely on beliefs in propositions that are indeterminate in truth value cannot rise to the status of determinate knowledge. Moreover, it is tempting to think that inference from indeterminately known propositions must itself result in indeterminate knowledge — even when the inferred proposition is determinate. Such temptations must, you guessed it, be resisted. It is true that in some cases one may be in a state of indeterminate knowledge towards a proposition that is determinately true. But in others, one may simply acquire determinate knowledge on the basis of inference from propositions of which one merely has indeterminate knowledge. Let me illustrate both possibilities.

that it will be sunny tomorrow, God cannot determinately know that it will be sunny. But God cannot determinately fail to know it either, for that would completely undermine omniscience. The alternative would be to describe God as indeterminately knowing each one of the relevant open historical possibilities. And the key difference is that, according to this picture, God would have to be understood as indeterminately *believing* each of those propositions.

[18] Of course adding the axiom is also open to the skeptic, in the same trivializing sense in which the skeptic accepts the lifting principles.

Start with a case in which it is unproblematic to describe an agent as having indeterminate knowledge in a determinate truth. As a background, note that, even if we deny any kind of reductive analysis, it is plausible that knowledge may be linked to various kinds of necessary conditions. The exact list of these conditions does not matter, but let us say for the sake of illustration that they include belief, safety and truth. The important observation is that indeterminacy may affect any one of these requirements.

It helps to reflect on the structure of a scenario in which it is indeterminate whether one's belief is safe. To construct this structure, suppose (i) I determinately believe both **A** and **C**, and in particular I believe **C** only on the basis of inference from **A** (ii) **A** is indeterminate and indeterminately known; (iii) **C** is determinately true, and (iv) for both beliefs it is indeterminate whether they are safe. This type of scenario illustrates that as a matter of principle, there is nothing surprising in the idea that one may indeterminately know a determinate truth. Moreover, the possibility that safety might be indeterminate is not entirely far fetched. After all, the safety constraint is spelled out in terms of a notion of modal closeness and, in the nearby counterfactuals literature, Stalnaker has famously argued that the closeness relation that matters to counterfactuals is sometimes indeterminate (Stalnaker, 1981).

Next consider some equally harmless scenarios in which one acquires items of determinate knowledge from states of indeterminate knowledge. Some authors have pointed to the possibility of acquiring knowledge on the basis of inference from false beliefs (Warfield, 2005; Fitelson, 2010). Here is a version of Warfield's core case: I look at my watch, it says that it is 3:45; on the basis of that, I form the belief that it is 3:45; from that belief, I infer that I am on time for my 4 PM appointment. However, unbeknownst to me, my watch is a minute fast, and it is . My belief that it is 3:45 is not knowledge, because it is not true, but plausibly my belief that I am on time for my appointment is. Cases of inferring a determinate truth from an indeterminate premise may have this same general structure. Imagine this scenario: I determinately believe and indeterminately know that it will rain; I infer from that belief that it will rain that I will need boots to be comfortable; however, that inferred belief is determinately true, because it is determinately true that it will either rain or snow tomorrow. The idea that this is possible is strictly less spooky than the already merely mildly spooky idea of inferring knowledge from falsehood.

9 The classical modal analysis

To conclude, let us explore a model theoretic analysis of the logical points I have made along the way. One of the most prominent models of indeterminacy, roughly codified in Barnes and Williams (2011) and applied to the open future in Barnes and Cameron (2009, 2011) is *precisificationist*. According to this model, indeterminacy in what world is actual is to be understood in terms of there being multiple, incompatible, and fully precise representations, each of which has equal claim to accurately represent our reality. The proposal is avowedly non-reductivist about the relevant sort of indeterminacy: "having equal claim" is just a matter of it being unsettled which fully precise representation is correct. One of the main points of emphasis for this view is that indeterminacy does not show up in the basic semantic analysis of the relevant fragment of language, which is otherwise endowed with a thoroughly classical logic.

It is possible to complement this proposal with a classical model theory for a multi-modal logic.[19] The key philosophical application for this model theory is that it will help us pin down some *non*-entailment relations which we have only been able to gesture at.

Recall that, given a normal modal logic for K and a normal modal logic for D, the fusion **LKD** of the logics is the smallest modal logic that extends both. Its model theoretic characterization, for our language, is provided by models of the following form:

Definition 1 *A* fusion model *is a structure $\langle W, R_K, R_D, v\rangle$ such that W is a set of worlds, R_K and R_D are respectively the accessibility relations for K and D, and v is a valuation function.*

Relative to fusion models we give the straightforward semantics

I. for each atomic sentence A, $\mathcal{M}, w \models \mathsf{A}$ iff $v_{\mathcal{M}}(\mathsf{A}, w) = 1$

II. for any sentences A, B:

- $\mathcal{M}, w \models \mathsf{A}\ \&\ \mathsf{B}$ iff $\mathcal{M}, w \models \mathsf{A}$ and $\mathcal{M}, w \models \mathsf{B}$

[19] By developing the classical model theory, I do not mean to rule out the possibility or interest of a non-classical analysis for the modals that retains classical logic for the non-modal part of the language. Indeed, I think there is promise in leveraging Humberstone's possibility semantics (Humberstone, 1983) to generate an alternate non-precisificationist model of the indeterminacy of the future. But this is a line of analysis that presents unique technical and conceptual problems and needs to be argued for and developed separately. Moreover, for the purposes of analyzing our principles, the classical approach will suffice.

– $\mathcal{M}, w \models \mathsf{A} \vee \mathsf{B}$ iff $\mathcal{M}, w \models \mathsf{A}$ or $\mathcal{M}, w \models \mathsf{B}$

– $\mathcal{M}, w \models \neg\mathsf{A}$ iff $\mathcal{M}, w \not\models \mathsf{A}$

III. for any A,

– $\mathcal{M}, w \models D\mathsf{A}$ iff for all u, $wR_D u$, $\mathcal{M}, u \models \mathsf{A}$

– $\mathcal{M}, w \models K\mathsf{A}$ iff for all u, $wR_K u$, $\mathcal{M}, u \models \mathsf{A}$

The logic generated by fusion models (given appropriate constraints on R_K and R_D) is the logic is generated by the union of the axioms of the individual logics (Carnielli and Coniglio, 2020, §4.1) . This is not the same as the union of the logics, since for instance $\neg DA \to \neg DKA$ is not a theorem of either logic but it is a theorem of the fusion.

In the course of the above discussion, I noted three principles that need to be added to the fusion, These can now be given precise model-theoretic treatment.

Fact 1 *The skeptical bridge ($\Diamond A \to \blacklozenge A$, or equivalently $\neg DA \to \neg KA$) is not valid over the class of all fusion frames. This principle is characterized are by the frame condition:*

$$\forall x, y(xR_D y \to xR_K y)$$

Fact 2 *The determinacy of knowledge ($KA \to DKA$) is not valid over the class of all fusion frames. It is characterized by the frame condition:*

$$\forall x, y((xR_D y) \to \forall z(yR_K z \to xR_K z))$$

Fact 3 *The unidirectionality of indeterminate knowledge ($IKA \ \& \ IKB \to \Diamond(A \ \& \ B)$) is not valid over the class of all fusion frames. It is characterized over these frames by the condition:*

$$\forall x, y, z(xR_D y \ \& \ xR_D z \to \exists w(yR_K w \ \& \ zR_K w \ \& \ xR_D w))$$

Proofs of these facts are routine checks with respect to the relevant class of frames.[20]

These simple technical results illuminate the structure that skeptical and non-skeptical stances impose on the relationship between epistemic and historical possibilities. Less formally, the skeptical bridge corresponds to the claim that historical possibilities are epistemic possibilities; the determinacy of knowledge

[20] At the time of this writing, these can be verified by computer by running the SQEMA algorithm for described in Conradie *et al.* (2006) and implemented at `https://store.fmi.uni-sofia.bg/fmi/logic/sqema/sqema_gwt_20180317_2/index.html`.

corresponds to the claim that any two historical possibilities must be related to the same epistemic possibilities. Finally, the unidirectionality of indeterminate knowledge corresponds to the claim that any two historical possibilities must access a common epistemic possibility.

Reflecting on the frame conditions in Facts 1-3 helps fully pin down the relationships between the principles. The skeptical bridge follows from the determinacy of knowledge only on the basis of reflexivity of R_K (just consider $y = z$ and notice that reflexivity yields yR_Ky). The converse entailment follows only on the controversial assumption that R_K is transitive.[21] This is the semantic correlate of the KK principle, which as we noted in footnote 8 can be used to establish this entailment. The frame condition for unidirectionality follows from the skeptical core as well, provided R_K is Euclidean. Of course it can also be imposed *independently* of the skeptical core principle, since it does not itself entail any of the skeptical principles.

Finally, the model theoretic analysis can also help illustrate the dynamics of indeterminate knowledge as discussed in section 5. Focus on the relation R_D that tracks the historical possibilities. The thought is that R_D is not a static relation but one that evolves over time.

Assume for simplicity that R_D is an equivalence relation, and thus that the cells it generates are partitions. As I mentioned, this is a controversial assumption in the indeterminacy literature, but seems plausible in the case of the open futurist position. So suppose that on Monday we start with a scenario like the one in Figure 3.

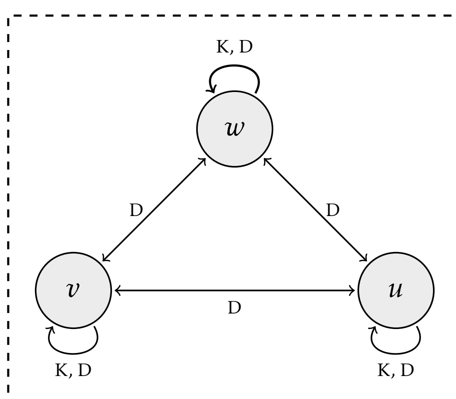

Figure 3: Monday's state

[21] By the skeptical core condition xR_Dy gets us xR_Ky; yR_Kz and transitivity we get xR_Kz. If transitivity fails, it is easy to provide a countermodel.

Suppose that A is true at v and u but false at w and consider the specific perspective of world v. We have $\mathcal{M}, v \models I$A and crucially $\mathcal{M}, v \models IK$A. (This is because $\mathcal{M}, v \models K$A while the historically possible world w has $\mathcal{M}, w \models \neg K$A.)

Imagine now that between Monday and Tuesday we learn the proposition corresponding to the set $\{v, u\}$ and consider Tuesday's perspective. Recall that, even though we are considering Tuesday's determinacy statuses, we are still interested in how those determinacy facts affects the determinacy status of Monday's knowledge states. The upshot of this is that the relation R_K is not allowed to change between the two states of the model, while we assume instead that the accessibility relation R_D has evolved. The historical possibilities at v now only include u in addition to v itself. For this reason we have $\mathcal{M}, v \models D$A and similarly $\mathcal{M}, v \models DK$A.

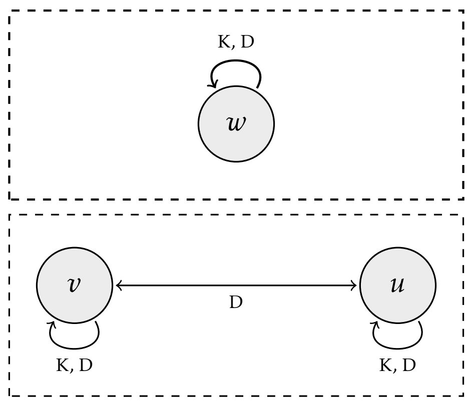

Figure 4: Tuesday state

This models the the sense in which on Tuesday it becomes determinate that on Monday one determinately knew **A**. The implied general picture is one on which for time-fixed K we consider a series of models with D-accessibility relations that are nested, in the sense that, as the series, progresses "later" partitions are included in "earlier" partitions.

Bibliography

Aronowitz, Sara and Lombrozo, Tania (2020), 'Learning through simulation', *Philosophers Imprint*.

Balcerak Jackson, Magdalena (2018), 'Justification by imagination', in Fiona Macpherson and Fabian Dorsch (eds.), 'Perceptual Imagination and Perceptual Memory', 209–226, Oxford University Press.

Barnes, Elizabeth and Cameron, Ross (2009), 'The open future: bivalence, determinism and ontology', *Philosophical Studies*, 146, 291–309.

Barnes, Elizabeth and Cameron, Ross (2011), 'Back to the open future', *Philosophical Perspectives*, 25(1), 1–26.

Barnes, Elizabeth and Williams, J.R.G. (2011), 'A theory of metaphysical indeterminacy', *Oxford studies in metaphysics*, 6, 103–148.

Barnett, David (2009), 'Is vagueness sui generis ?', *Australasian Journal of Philosophy*, 87(1), 5–34.

Belnap, Nuel and Green, Mitchell (1994), 'Indeterminism and the thin red line', *Philosophical perspectives*, 3, 365–388.

Belnap, Nuel, Perloff, Michael and Xu, Ming (2001), *Facing the Future*, Oxford University Press.

Besson, Corine and Hattiangadi, Anandi (2014), 'The open future, bivalence and assertion', *Philosophical Studies*, 167(2), 251–271.

Bird, Alexander (2007), 'Justified judging', *Philosophy and Phenomenological Research*, 74(1), 81–110.

Borghini, Andrea and Torrengo, Giuliano (2013), 'The metaphysics of the thin red line', in Fabrice Correia and Andrea Iacona (eds.), 'Around the Tree', 105–125, Springer.

Caie, Michael (2012), 'Belief and indeterminacy', *Philosophical Review*, 121(1), 1–54.

Cameron, Ross (2015), *The Moving Spotlight*, Oxford University Press.

Cariani, Fabrizio (2020), 'On predicting', *Ergo*, 7(11), 339–361.

Cariani, Fabrizio (2021), *The Modal Future*, Cambridge University Press.

Cariani, Fabrizio and Santorio, Paolo (2018), '*Will* done better: Selection semantics, future credence, and indeterminacy', *Mind*, 127(505), 129–165.

Carnielli, Walter and Coniglio, Marcelo Esteban (2020), 'Combining Logics', in Edward N. Zalta (ed.), 'The Stanford Encyclopedia of Philosophy', Metaphysics Research Lab, Stanford University, fall 2020 edition.

Carter, Sam and Goldstein, Simon (2020), 'The normality of error', *Philosophical Studies*, 1–25.

Christensen, David (2004), *Putting Logic in Its Place*, Oxford University Press.

Conradie, Willem, Goranko, Valentin and Vakarelov, Dimiter (2006), 'Algorithmic correspondence and completeness in modal logic. i. the core algorithm sqema', *arXiv preprint cs/0602024*.

Correia, Fabrice and Iacona, Andrea (2013), *Around the tree*, Springer.

Dorr, Cian (2003), 'Vagueness without ignorance', *Philosophical Perspectives*, 17(1), 83–113.

Dorr, Cian, Goodman, Jeremy and Hawthorne, John (2014), 'Knowing against the odds', *Philosophical Studies*, 170(2), 277–287.

Field, Hartry (2000), 'Indeterminacy, degree of belief, and excluded middle', *Noûs*, 34(1), 1–30.

Fitelson, Branden (2010), 'Strengthening the case for knowledge from falsehood', *Analysis*, 70(4), 666–669.

Foley, Richard (1979), 'Justified inconsistent beliefs', *American Philosophical Quarterly*, 16(4), 247–257.

Goldman, Alvin I. (1967), 'A causal theory of knowing', *The journal of Philosophy*, 64(12), 357–372.

Goodman, Jeremy and Salow, Bernhard (2018), 'Taking a chance on kk', *Philosophical Studies*, 175(1), 183–196.

Goodman, Jeremy and Salow, Bernhard (ms.), 'Epistemology normalized', USC and Oxford University.

Grice, H. Paul (1975), 'Logic and conversation', in 'Studies in the Ways of Words (1989)', 22–40, Harvard University Press.

Hawthorne, John and Lasonen-Aarnio, Maria (2009), 'Knowledge and objective chance', in Patrick Greenough and Duncan Pritchard (eds.), 'Williamson on Knowledge', 92–108, Oxford University Press.

Hirsch, Eli (2006), 'Rashi's view of the open future: Indeterminateness and bivalence', *Oxford Studies in Metaphysics*, 2, 111–135.

Humberstone, Lloyd (1983), 'The background of circumstances', *Pacific Philosophical Quarterly*, 64, 19–34.

Iacona, Andrea (2013), 'Timeless truth', in Fabrice Correia and Andrea Iacona (eds.), 'Around the Tree', 29–45, Springer.

Ichikawa, Jonathan Jenkins (2014), 'Justification is potential knowledge', *Canadian Journal of Philosophy*, 44(2), 184–206.

Kahneman, Daniel and Tversky, Amos (1982), 'The simulation heuristic', in Daniel Kahneman, Stewart Paul Slovic and Amos Tversky (eds.), 'Judgment under uncertainty: Heuristics and biases', 201–208.

Kurucz, Agi (2007), 'Combining modal logics', *Handbook of modal logic*, 3, 869–924.

Littlejohn, Clayton and Dutant, Julien (2019), 'Justification, knowledge, and normality', *Philosophical Studies*, 1–17.

Łukasiewicz, Jan (1970), 'On determinism', in Ł. Borkowski (ed.), 'Jan Łukasiewicz: Selected Works', .

MacFarlane, John (2003), 'Future contingents and relative truth', *Philosophical Quarterly*, 53, 321–336.

MacFarlane, John (2014), *Assessment Sensitivity*, Oxford University Press.

Malpass, Alex and Wawer, Jacek (2012), 'A future for the thin red line', *Synthese*, 188, 117–142.

Neta, Ram (2018), 'Evidence, coherence and epistemic akrasia', *Episteme*, 15(3), 313–328.

Pollock, John L and Cruz, Joseph (1999), *Contemporary theories of knowledge*, volume 35, Rowman & Littlefield.

Smith, Martin (2016), *Between probability and certainty: What justifies belief*, Oxford University Press.

Staffel, Julia (2019), 'How do beliefs simplify reasoning?', *Noûs*, 53(4), 937–962.

Stalnaker, Robert (1981), 'A defense of conditional excluded middle', in 'Ifs', 87–104, Springer.

Stalnaker, Robert (2006), 'On logics of knowledge and belief', *Philosophical studies*, 128(1), 169–199.

Stalnaker, Robert C. (2019), *Knowledge and conditionals: essays on the structure of inquiry*, Oxford University Press.

Thomason, Richmond H. (1970), 'Indeterminist time and truth-value gaps', *Theoria*, 18(3), 264–281.

Todd, Patrick and Rabern, Brian (2021), 'Future contingents and the logic of temporal omniscience', *Noûs*, 55(1), 102–127.

Torre, Stephan (2011), 'The open future', *Philosophy Compass*, 6(5), 360–373.

Torrengo, Giuliano and Iaquinto, Samuele (2020), 'The invisible thin red line', *Pacific Philosophical Quarterly*, 101, 354–382.

Varzi, Achille C (2007), 'Supervaluationism and its logics', *Mind*, 116(463), 633–676.

Warfield, Ted A (2005), 'Knowledge from falsehood', *Philosophical Perspectives*, 19, 405–416.

Williams, J Robert G (2012), 'Indeterminacy and normative silence', *Analysis*, 72(2), 217–225.

Williams, J Robert G (2014), 'Nonclassical minds and indeterminate survival', *Philosophical Review*, 123(4), 379–428.

Williamson, Timothy (1994), *Vagueness*, Routledge.

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

Williamson, Timothy (2008), *The philosophy of philosophy*, John Wiley & Sons.

Williamson, Timothy (2009), 'Reply to Hawthorne and Lasonen-Aarnio', in Patrick Greenough and Duncan Pritchard (eds.), 'Williamson on Knowledge', 313–329, Oxford University Press.

Worsnip, Alex (2016), 'Belief, credence, and the preface paradox', *Australasian Journal of Philosophy*, 94(3), 549–562.