Simple Semantics for Logics of Indeterminate Epistemic Closure

COLIN R. CARET

Abstract: According to Jago (2014a), logical omniscience is really part of a deeper paradox. Jago develops an epistemic logic with principles of indeterminate closure to solve this paradox, but his official semantics is difficult to navigate, it is motivated in part by substantive metaphysics, and the logic is not axiomatized. In this paper, I simplify this epistemic logic by adapting the hyperintensional semantic framework of Sedlár (2021). My first goal is metaphysical neutrality. The solution to the epistemic paradox should not require appeal to a metaphysics of truth-makers, situations, or impossible worlds, by contrast with Jago's official semantics. My second goal is to elaborate on the proof theory. I show how to axiomatize a family of logics with principles of indeterminate epistemic closure.

Keywords: knowledge, closure, paradox, omniscience, hyperintensionality

1 Introduction

On the face of it, logical omniscience is just an artifact of certain semantic techniques (see, e.g. Hintikka 1962). In order to avoid this problem, we just need to find a semantics for knowledge that is not so crude.

If knowledge is not closed under all logical consequences, then a natural follow-up question is whether it is closed under some *restricted class* of logical consequences.[1] We might be tempted to think about this as follows. Some logical consequences are obvious to all epistemic subjects, in virtue of their capacity for rational thought. Call them *trivial* consequences. The inference from "φ and ψ" to its conjunct φ could be a plausible example of triviality. Knowing the individual conjunct seems like it is part of knowing the whole conjunction.[2] It just comes along for free. Although it is difficult to precisely define triviality, the *primitive rules* of proof theory seem to fall into this category because they are the simplest deductive inferences.

[1] Cf. Duc (1995) and Jago (2006) on 'logical ignorance'.

[2] Yablo (2014) and Hawke, Özgün, and Berto (2020) defend this claim.

According to Jago (2014a), however, this reveals that logical omniscience is actually part of a deeper problem. On the one hand, we want to deny that knowledge is closed under all logical consequences. On the other hand, we want to assert that knowledge is 'trivially closed', i.e. it is closed under consequences that are derivable by a single use of any primitive rule. The problem is that this is paradoxical: any logical consequence is derivable by a chain of primitive rules, so trivial closure implies full closure.

Jago's solution is a theory of Indeterminate Epistemic Closure (IEC), viz. the view that knowledge closure is only partially determinate and trivial consequences are borderline cases where this determinacy breaks down. On this view, the following principle should hold in an epistemic logic enhanced with operator Δ to be read 'it is determinate that…' (Jago, 2014a, p.251).

(IEC) If $\varphi_1, \ldots, \varphi_n \vDash \psi$ is trivial, $\Delta K\varphi_1, \ldots, \Delta K\varphi_n \vDash \neg\Delta\neg K\psi$

We can gloss IEC as follows: with respect to any trivial inference, it is impossible for 'the break down of knowledge closure' to be determinate. More carefully, it is impossible for a subject to determinately know the premises and determinately fail to know the conclusion of a trivial inference. This notion of determinacy is supposed to be something like the notion used in discussions of vagueness.[3] For Jago, the philosophical importance of this concept largely comes out as a norm of assertion: "…we can never rationally assert that such-and-such is an epistemic oversight…Such cases are always indeterminate cases and as such do not rationally support assertions about them in the way that clear cases do." (Jago, 2014a, p.19)

On this view, when knowledge closure *does* break down over a trivial inference step, it is unassertible *that* it breaks down there. This yields a diagnosis of the paradox as follows. Assuming that the role of triviality is captured by IEC, we should recognize that for any trivial inference, it is indeterminate that knowledge closure breaks down at that step. We might then assert: closure does not break down at that step! This is a mistake, but an easy one to make because this can sound very similar to IEC itself.[4] Once we fall prey to this mistake, however, it compels us to think that knowledge is trivially closed and thereby draws us into the paradox.[5]

[3]Where it often used to define borderline cases, i.e. where it is *borderline* φ if, and only if, it is neither determinately φ nor determinately $\neg\varphi$.

[4]The mistake consists in sliding from a claim of the form $\neg\Delta\chi$ that says there is a certain *failure of assertibility*, to a claim of the form $\Delta\neg\chi$ which says there is *an assertible failure*.

[5]Another option to consider is a non-classical approach to K, where the degree of truth of $K\varphi$ is low but non-zero for any conclusion φ of a trivial inference. This would model

In this paper, I will explore a family of IEC logics with two main goals. My first goal is metaphysical neutrality. The problem and solution described above made perfectly good sense without any metaphysical assumptions. In particular, it did not seem to involve any mention of entities such as impossible worlds (as used in Jago's semantics). I will show that a simple semantics is available without appeal to substantive metaphysical categories. This is not because I believe that impossible worlds are problematic in principle, but simply because it seems to me that the formal solution to the paradox is conceptually independent of our metaphysical commitments.

My second goal is to elaborate on the proof theory. Although some properties of the knowledge operator and determinacy operator are clear from Jago's exposition, he has not given a systematic proof theory for IEC logics. I will take a step in that direction. To achieve these ends, I will leverage some recent work on hyperintensional semantics by Sedlár (2021). In the next section, I present an overview of this framework.

2 Sphere models for determinate knowledge

What are the target properties of this epistemic logic? I will start with some negative properties. For one thing, knowledge ought not be closed under replacement of necessary equivalents, i.e. it is a hyperintensional context. Similar intuitions to those that lead us to reject omniscience in the first place should also lead us to reject closure under logical equivalence. For another thing, determinate truth ought to be stronger than 'mere truth'.

- $\varphi \leftrightarrow \psi \nvDash K\varphi \leftrightarrow K\psi$
- $\varphi \nvDash \Delta\varphi$[6]

As for positive properties: both operators are factive, and the determinacy operator is a normal modality (e.g. tautologies are determinately true).

- $K\varphi \vDash \varphi$
- $\Delta\varphi \vDash \varphi$

indeterminacy from the metalanguage, instead of adding a marker for determinacy to the classical object-language. Thanks to the editor for this thought.

[6]This is one respect in which the present treatment of determinacy differs from that of supervaluationist logics (see, e.g. Fine 1975). There is also a subtle conceptual difference in that we have in mind an objective rather than subjective kind of determinacy.

- $\vDash \varphi$ implies $\vDash \Delta\varphi$

In the semantics below, I will follow Jago's lead and treat the determinacy operator as a strong normal modality like the S5 necessity operator, but in principle it could be understood as some kind of weaker operator.

To articulate an elegant semantics that combines these two operators, I will adapt the hyperintensional semantics of Sedlár (2021). Every sentence has both a fine-grained content (FGC) and a coarse-grained content (CGC) in this semantic framework. These divide the traditional work of a *proposition*. The CGC of a sentence is simply a set of possible worlds or a 'truth set', which can be used to define logical properties like validity.

The FGCs of sentences, on the other hand, are genuine primitives in the sense that they are *not* inter-definable with other components of the framework. As a type of content, a FGC partitions worlds into those where it is true and those where it is not, but it is distinct from this corresponding 'truth set'. This allows for a two-level, composite analysis of the contents of sentences, whereby sentences are first assigned a FGC, and this in turn determines the 'truth set' or CGC of that sentence.

This is a generalization of neighborhood semantics. The *hyperintension* function H maps each sentence to a FGC, which the *intension* function I then maps to a CGC, and their composition satisfies classical operations on 'truth sets' such as $I(H(\neg\varphi)) = W \setminus I(H(\varphi))$. The knowledge operator is then interpreted along the lines of a modal operator from neighborhood semantics, i.e. there is a neighborhood function that has the job of assigning a set of 'known contents' to the designated agent. However, this is not a set of CGCs as in traditional neighborhood semantics, but a set of FGCs.

To extend this framework with a (normal) determinacy operator, we note that the access relation R of a Kripke model could always be replaced with a function $S(w) = \{x : wRx\}$ that outputs a *sphere* of alternatives to w. We assume that some truths are not determinate. This is represented by the existence of worlds in the sphere that disagree with each other. Determinate truths are those that hold the same in all worlds throughout this sphere.

A case of the IEC principle can then be captured by roughly the following modeling condition: whenever the FGC of φ is in the knowledge set of all worlds in the sphere of w, the FGC of ψ is in the knowledge set of some world in the sphere of w. This ensures that the agent determinately knows φ only if it is *indeterminate* that they fail to know ψ. Does this mean that this semantic method requires us to define the class of trivial consequences, once and for all? Fortunately not, as I will explain later.

Let me first give the basic semantics for the following signature.

$$\mid p_i \mid \neg \mid \bot \mid \ \rightarrow \ \mid K \mid \Delta \mid$$

Definition 1 (Sphere Models) *A sphere model $M = \langle W, C, H, I, N, S \rangle$ is a six element structure such that…*

- *$W \neq \emptyset$ is a set of possible worlds*
- *$C \neq \emptyset$ is a set of fine-grained contents or FGCs*
- *H : Sent $\rightarrow C$ is a hyperintension function*
- *$I : C \rightarrow \wp(W)$ is an intension function*
- *$N : W \rightarrow \wp(C)$ is a neighborhood function, such that*
  - *If $c \in N(w)$, then $w \in I(c)$.*
- *$S : W \rightarrow \wp(W)$ is a sphere function, such that*
  - *Spheres are centered and jointly partition W.*

In addition, the coarse-grained content or CGC of φ in model M, given by the function $[\![\varphi]\!]^M = I(H(\varphi))$, must satisfy the following constraints.

- $[\![p]\!]^M = I(H(p))$ *for atomic* p
- $[\![\bot]\!]^M = \emptyset$
- $[\![\neg\varphi]\!]^M = W \setminus [\![\varphi]\!]^M$
- $[\![\varphi \rightarrow \psi]\!]^M = (W \setminus [\![\varphi]\!]^M) \cup [\![\psi]\!]^M$
- $[\![K\varphi]\!]^M = \{x \in W : H(\varphi) \in N(x)\}$
- $[\![\Delta\varphi]\!]^M = \{x \in W : S(x) \subseteq [\![\varphi]\!]^M\}$

The functions H and I can be naturally extended to sets of sentences and their contents, e.g. $H(\Gamma) = \{H(\psi) : \psi \in \Gamma\}$. For any given world $w \in W$ of a model, I refer to $N(w)$ as the *knowledge set* of that world. This is a framework for single-agent knowledge claims, but it can easily be extended to a multi-agent setting by parameterizing K and N to a set of agent names. Since this does not illuminate anything particularly interesting about the solution to the epistemic paradox, I leave aside such details.

The truth-conditions of a sentence are reflected in their CGCs, so we have that $K\varphi$ is true iff the FGC of φ is in the knowledge set according to w, and $\Delta\varphi$ is true iff φ is true throughout the sphere of w. The modeling condition on the neighborhood function, which coordinates known contents in N with intensions assigned by I, ensures factivity.

Here is a partial representation of a sphere model M to illustrate the main ideas of the semantics. In this model we have a four element Boolean algebra over a set of contents $C = \{c_1, c_2, c_3, c_4\}$ and we have a set of two worlds $W = \{w_1, w_2\}$ with 'universal' spheres $S(w_1) = W = S(w_2)$.

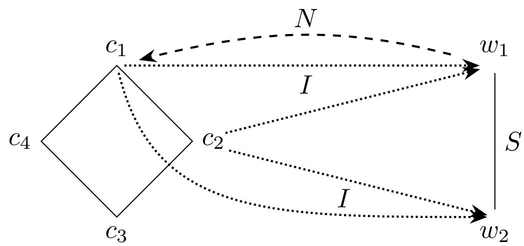

In a language of two atoms, we will let $H(p) = c_1$ and $H(q) = c_2$ and $H(\neg p) = c_3$ and $H(\neg q) = c_4$, with conjunctions and disjunctions assigned meets and joins of contents. Let $H(Kp) = c_2$ and $H(Kq) = c_3$ and $H(\Delta Kp) = c_3$. So, the top element is the FGC of p, the right is the FGC of q and Kp, and the bottom is the FGC of Kq and ΔKp.

The intension function maps $I(c_1) = W = I(c_2)$ as indicated in the diagram, but also $I(c_3) = \emptyset = I(c_4)$ which is not drawn. Finally, the neighborhood function maps $N(w_1) = \{c_1\}$ as indicated in the diagram, but also $N(w_2) = \emptyset$ which is not drawn. At world w_1, both p and q are true, but they are epistemically distinct because only p is known, furthermore the knowledge of p is not a determinate fact at w_1 because there is another world in its sphere w_2 where p is not known. Formally, we have:

- $[\![p]\!]^M = W$
- $[\![q]\!]^M = \{w_1\}$
- $[\![Kp]\!]^M \{w_1\}$
- $[\![Kq]\!]^M = \emptyset$
- $[\![\Delta Kq]\!]^M = \emptyset$

In addition to reading them as truth-conditions, we also use the CGCs of sentences to define logical consequence as (local) truth-preservation.

Definition 2 ($\mathbb{C}$-Consequence) *The consequence relation of a class $\mathbb{C}$ of sphere models is defined by: $\Gamma \vDash_{\mathbb{C}} \varphi$ iff $[\![\Gamma]\!]^M \subseteq [\![\varphi]\!]^M$ in all $M \in \mathbb{C}$.*

This achieves many of our desiderata. On this semantics, knowledge is not closed under replacement of necessary equivalents. Since logical equivalence just consists in $[\![\varphi]\!]^M = [\![\psi]\!]^M$ in all models, we can model equivalent sentences as having distinct FGCs, so one can be known without the other. Notably, this 'fine-graining' is achieved without appeal to the metaphysics of truth-makers, situations, or impossible worlds. In addition, the knowledge and determinacy operators are factive and the determinacy operator is 'S5-like'. The modeling conditions on the sphere function ensure this because '...is in the sphere of...' is an equivalence relation.[7] These facts about sphere semantics are recorded below.

Remark 1 In the class $\mathbb{C}$ of all sphere models we have:

- $\varphi \leftrightarrow \psi \nvDash_{\mathbb{C}} K\varphi \leftrightarrow K\psi$ **(K-Hyperintensionality)**
- $\varphi \nvDash_{\mathbb{C}} \Delta\varphi$ **(Δ-Strength)**
- $K\varphi \vDash_{\mathbb{C}} \varphi$ **(K-Factivity)**
- $\Delta\varphi \vDash_{\mathbb{C}} \varphi$ **(Δ-Factivity)**
- $\neg\Delta\neg\varphi \vDash_{\mathbb{C}} \Delta\neg\Delta\neg\varphi$ **(Δ-Euclidean)**
- $\Delta(\varphi \rightarrow \psi) \vDash_{\mathbb{C}} \Delta\varphi \rightarrow \Delta\psi$ **(Δ-Distribution)**
- $\vDash_{\mathbb{C}} \varphi$ implies $\vDash \Delta\varphi$ **(Δ-Necessitation)**

What this does *not* yet achieve is the coordination of determinacy and knowledge operators that is required to validate the IEC principle. In the next section, I will explain how we can achieve this without actually defining the class of trivial consequences in advance, once and for all. The vagueness of triviality regulates this concept in a way that can be formally modeled.

[7] By the centering requirement, we always have $w \in S(w)$, and since a sphere is part of a partition, if $x \in S(w)$, then x and w simply have the same sphere $S(x) = S(w)$.

3 Respecting and regulating triviality

As a warm-up, I show how to validate an arbitrary case of IEC, assuming that we have identified a particular inference form as trivial. An inference form is a SET-FML pair. Suppose that the inference form $\langle \Gamma, \varphi \rangle$ is trivial. Call the following condition the Γ-φ *respect schema* (schematic over w).

(Γ-φ-RS) $H(\Gamma) \nsubseteq \bigcap\{N(x) : x \in S(w)\}$ or $H(\varphi) \in \bigcup\{N(x) : x \in S(w)\}$

If this holds at a given world w where it is a determinate truth that all the Γs are known, then it is indeterminate, at w, that knowledge of φ fails. This will be spelled out more carefully below. The upshot is that when this condition holds globally in a class of models, they validate the $\langle \Gamma, \varphi \rangle$ case of IEC.

Theorem 1 (The Γ-φ Respect Theorem) *For any given inference form $\langle \Gamma, \psi \rangle$, define a class $\mathbb{C}$ of sphere models as follows: $M \in \mathbb{C}$ iff* (Γ-φ-*RS*) *holds at all $w \in W$ of M. Then $\{\Delta K\psi : \psi \in \Gamma\} \vDash_{\mathbb{C}} \neg\Delta\neg K\varphi$.*

Proof. Let $M \in \mathbb{C}$ and $M = \langle W, C, H, I, N, S \rangle$. Let $w \in W$ and, for all $\psi \in \Gamma$, let $w \in [\![\Delta K\psi]\!]^M$. Then, for all $\psi \in \Gamma$ and all $x \in S(w)$, we have $x \in [\![K\psi]\!]^M$, equivalently, we have $H(\psi) \in N(x)$. So, for all $x \in S(w)$, we have $H(\Gamma) \subseteq N(x)$. In that case, $H(\Gamma) \subseteq \bigcap\{N(x) : x \in S(w)\}$ and so the condition (Γ-φ-RS) implies that $H(\varphi) \in \bigcup\{N(x) : x \in S(w)$ holds. It follows that $H(\varphi) \in N(x)$ for some $x \in S(w)$. So, $x \in [\![K\varphi]\!]^M$ for some $x \in S(w)$. Thus, $S(w) \nsubseteq [\![\neg K\varphi]\!]^M$ and so $w \in [\![\neg\Delta\neg K\varphi]\!]^M$. □

The question, then, is whether we can say anything more interesting about the logical role of triviality. Hypothetically, one could argue that we have an adequate solution to the paradox so long as we understand what it looks like to establish a modeling condition (Γ-φ-RS) for *whatever* inference forms $\langle \Gamma, \psi \rangle$ are trivial. Perhaps the question of which inference forms are genuinely trivial is only suited to informal philosophical debate.

This attitude, however, is unsatisfying if we want to develop the proof theory of IEC logics. If we leave 'trivial consequence' completely unsettled, it makes no sense to even try to axiomatize the relevant cases of the IEC principle (what are the relevant cases?). For this reason, I will attempt a more formal treatment of triviality. The trouble is that triviality is a vague concept. When applied to deductive inferences, 'trivial' connotes something like 'undeniably obvious'. There may be paradigm cases of this phenomenon, but there are also many contestable, borderline cases.

At the outset, I suggested that primitive rules represent paradigm cases of triviality. Jago (2014a, p.11 and p.163) also endorses this claim and argues that it underwrites a *ranking* of candidate extensions of triviality. On the one hand, this implies that proof-theoretic factors constrain the concept of triviality. On the other hand, this claim does not take any official stand on the correct extension of the concept. I will show that this ranking approach does all of the work we need to axiomatize IEC logics.

In more detail, the ranking approach works as follows. Derivations are recursively defined. Longer derivations are formed by extending shorter derivations with a single application of one of the primitive rules. The *rank* of a consequence is the length of its shortest derivation. For each $n \in \mathbb{N}$, this gives us a precise candidate extension for the concept of triviality, by collecting all of the valid inference forms of lower ranks.[8]

The idea is that the proof theory of the base logic (without K and Δ) *grades* logical consequences. Semantics alone is not suited to this task because it only classifies inference forms as valid or invalid, without further distinctions. The gradation can be seen as a proxy for relative opacity: higher ranking consequences are those that are harder to recognize, so they are less fitting candidates for triviality.[9] What is interesting about this analysis is that it implies nothing about the correct extension of the concept of triviality, it merely tells us that there is a smooth transition from one candidate extension to another (which reflects the vagueness of triviality).

Nonetheless, this analysis provides a lot of interesting formal structure to work with. Since the primitive rules of the base logic establish the ranking, it can be implemented relative to different choices of base logic, so at an abstract level, this analysis is quite neutral about 'what is really trivial'. It does, however, imply that there are constraints on triviality once a base logic is fixed in the background. To assert that a high-ranking consequence is trivial implies that any lower-ranking consequence is also trivial.

For my purposes, the most useful aspect of this analysis is that it makes axiomatization possible. In the next section, I will implement the ranking approach relative to a Hilbert system and apply that ranking to give a full

[8]This strategy relates epistemic logic to the topic of proof complexity. Duc (1995) was the first person to mention this idea. Artemov and Kuznets (2014) use this idea to give a treatment of omniscience with operators $\Box_n$ meaning '. . . is known after n proof steps'.

[9]To say that one candidate is less fitting than another does not imply that it is unfit as such. All of these sets can be seen as genuine candidates in the sense that they could conceivably demarcate the extension of the concept of triviality.

semantics for a family of epistemic logics. This is best understood as one fully worked example of the analysis presented above.

4 A family of IEC logics

Consider a list of *potential* principles relating knowledge closure to certain valid inference forms of Classical Propositional Logic (CPL). We might ask ourselves: is knowledge of a tautology such as Excluded Middle indeterminate, or what about knowledge of a conclusion drawn by Modus Ponens? (we are asking about whether such consequences are trivial or not)

- $\models \neg\Delta\neg K(\varphi \vee \neg\varphi)$
- $\Delta K\neg\neg\varphi \models \neg\Delta\neg K\varphi$
- $\Delta K(\varphi \to \psi), \Delta K\varphi \models \neg\Delta\neg K\psi$
- $\Delta K(\varphi \vee \psi), \Delta K\neg\varphi \models \neg\Delta\neg K\psi$
- $\Delta K(\varphi \to \psi), \Delta K(\psi \to \chi) \models \neg\Delta\neg K(\varphi \to \chi)$

According to IEC theory, the validity of such inferences is always relativized to a candidate extension of triviality. Each principle holds relative to some candidates for triviality, but none hold absolutely.

We can make this precise. To do so, I will use a Hilbert system for CPL defined over the following signature (with other connectives definable).

$$| \; p_i \; | \; \neg \; | \; \bot \; | \; \to \; |$$

In order to implement the ranking approach, we want to think of derivations as the products of rules. Hilbert systems also usually have axioms, but we can capture this by stating a rule with no conditions, (R2) below, that effectively says that we can freely extend any derivation with any axiom. In the first instance, any rule-generated extension of the empty sequence of sentences counts as a derivation. Here are the axioms.

Definition 3 (Axioms of CPL) *The set of* axioms AX_{CPL} *of the background logic, CPL, is the set of all instances of the following schemata.*

(A1) $\varphi \to (\psi \to \varphi)$

(A2) $(\varphi \to (\psi \to \chi)) \to ((\varphi \to \psi) \to (\varphi \to \chi))$

(A3) $(\neg\varphi \to \neg\psi) \to (\psi \to \varphi)$

We use the axioms to define the primitive rules. The format of the rules, below, may look unusual, but they just re-write the usual defining conditions of a Hilbert derivation. In this formulation, primitive rules define admissible extensions of derivations (sequences of sentences).

Definition 4 (Primitive Rules of CPL) *A* derivation from Γ *in CPL is a finite (possibly empty) sequence of sentences $\psi_1, \ldots, \psi_n$. N.B. in the rules below, δ is used as a variable for such sequences.*

(R1) *If δ is a derivation from Γ and $\varphi \in \Gamma$, then δ, φ is a derivation from Γ.*

(R2) *If δ is a derivation from Γ and $\varphi \in AX_{CPL}$, then δ, φ is a derivation from Γ.*

(R3) *If δ is a derivation from Γ and there are members of this sequence of the form ψ and $\psi \to \varphi$, then δ, φ is a derivation from Γ.*

(R1) is the rule of assumptions, (R2) is the rule of free use of axioms, and of course (R3) is Modus Ponens. The set of all derivations from Γ is the smallest set of sequences closed under (R1)-(R3).

For the definition of derivability *per se*, we eliminate the empty sequence. Thus, the simplest inferences that are defined as derivable *per se* are one step derivations of assumptions or axioms.

Definition 5 (Derivability in CPL) *φ is derivable from Γ in CPL, written $\Gamma \vdash_{CPL} \varphi$, iff there is a non-empty derivation from Γ ending in φ, i.e. a sequence of sentences $\psi_1, \ldots, \psi_n$ that satisfies the criteria of Def. 4 such that $\psi_n = \varphi$. This sequence* witnesses *that $\Gamma \vdash_{CPL} \varphi$.*

This is the familiar derivability relation for CPL, re-written in a slightly unusual format. It is, however, worth making this explicit in order to clarify how this proof theory ranks the logical consequences of CPL. If the primitive rules (R1)-(R3) above, represent paradigm cases of triviality, then a specific ranking of sets of CPL consequences follows directly.

Definition 6 (Rank of CPL Consequences) *By the completeness theorem for CPL, if $\Gamma \vDash_{CPL} \varphi$, then there is a non-empty set of sequences of sentences $W(\Gamma, \varphi) = \{\delta : \delta$ witnesses that $\Gamma \vdash_{CPL} \varphi\}$. For each $\delta \in W(\Gamma, \varphi)$, let its length $\text{len}(\delta)$ be equal to the number of sentences in the sequence. The* rank *of this consequence is $\#(\Gamma \vDash_{CPL} \varphi) = \min\{\text{len}(\delta) : \delta \in W(\Gamma, \varphi)\}$.*

This ranking is what we care about. In a moment, I will point out how to use this ranking in an interesting way to generate IEC logics. First, here is a list of some examples to illustrate how various logical consequences of CPL are ranked by this system. Remember, this is meant to model relative opacity: higher ranks are less fitting candidates for triviality.

Remark 2 Ranking some logical consequences of CPL.

- $\#(\Gamma \vDash_{\text{CPL}} \varphi) = 1$ when $\varphi \in \Gamma$. **(Assumptions)**
- $\#(\Gamma \vDash_{\text{CPL}} \varphi \to (\psi \to \varphi)) = 1$ **(Axioms)**
- $\#(\Gamma, \varphi, \varphi \to \psi \vDash_{\text{CPL}} \psi) = 3$ **(Modus Ponens)**
- $\#(\Gamma, \neg\varphi \vDash_{\text{CPL}} \varphi \to \bot) = 5$ **(Negativity)**
- $\#(\Gamma, \varphi \to \psi, \psi \to \chi \vDash_{\text{CPL}} \varphi \to \chi) = 7$ **(Transitivity)**
- $\#(\Gamma, \varphi \to \psi \vDash_{\text{CPL}} \varphi \to (\psi \wedge \varphi)) = 9$ **(Pooling)**

Candidate extensions are sets of inference forms. The nth candidate extension for triviality collects all of SET-FML pairs corresponding to logical consequences of rank up to n. This candidate is called n-Triviality.

Definition 7 (n-Triviality) *$\mathbb{T}_n = \{\langle\Gamma, \varphi\rangle : \#(\Gamma \vDash_{CPL} \varphi) \leq n\}$ is the collection of all inference forms classified as 'trivial' by the nth candidate extension for triviality. I will write $\mathbb{T}_n\langle\Gamma, \varphi\rangle$ to mean $\langle\Gamma, \varphi\rangle \in \mathbb{T}_n$.*

Lemma 1 (Properties of n-Triviality) *The definition of $\mathbb{T}_n$ has all of the properties below, which leads Jago (2014b, p.1165) to refer to n-Triviality as '(a kind of) consequence' relation.*[10]

- *If $\mathbb{T}_n\langle\Gamma, \varphi\rangle$, then $\Gamma \vDash_{CPL} \varphi$* ***(Classicality)***
- *$\mathbb{T}_n\langle\varphi, \varphi\rangle$* ***(Reflexivity)***
- *If $\mathbb{T}_n\langle\Gamma, \varphi\rangle$, then $\mathbb{T}_n\langle\Gamma \cup \Sigma, \varphi\rangle$* ***(Monotonicity)***
- *If $\mathbb{T}_n\langle\Gamma, \varphi\rangle$, then $\mathbb{T}_{n+1}\langle\Gamma, \varphi\rangle$ (but not vice versa)* ***(Heredity)***
- *$\mathbb{T}_n\langle\Gamma, \varphi\rangle$ and $\mathbb{T}_n\langle\Gamma \cup \{\varphi\}, \psi\rangle$ do* not *imply $\mathbb{T}_n\langle\Gamma, \psi\rangle$* ***('Cut' Fails)***

[10] This may seem odd since 'Cut' fails, but recent work on substructural logics makes this remark more reasonable (see, e.g., Cobreros, Egré, Ripley, & van Rooij, 2012).

Proof. Straightforward. □

All of these properties are important, but Heredity in particular makes clear how this approach models the vagueness of triviality. It shows that, on this definition, there is a smooth transition from one candidate extension to another, which is why there are borderline cases of triviality.

Before returning to the full epistemic logic, we note one more useful fact. If $\mathbb{T}_n\langle\Gamma,\varphi\rangle$, then from Classicality, and the compactness theorem and deduction theorem, it follows that there is a finite set of assumptions $\{\psi_1,\ldots,\psi_m\}\subseteq\Gamma$ such that $\vdash_{\text{CPL}}\ \psi_1\rightarrow(\ldots\rightarrow(\psi_m\rightarrow\varphi))$. Call this associated, derivable, nested implication sentence a *reductive implication*.

Definition 8 (Reductive Implications) *For each candidate extension $\mathbb{T}_n$, let RI_n be the set of all of its associated reductive implications.*

Note that an n-Trivial inference form can have more than one reductive implication. This redundancy is not important. We only need to know that there is at least one reductive implication for each n-Trivial inference.

We saw that if (Γ-φ-RS) holds at a world of a sphere model, then at that world, it is indeterminate whether knowledge is always closed under the inference from Γ to φ. The semantics for an IEC logic ought to *globally* respect every inference form that is considered to be trivial. I will use the precise candidates for triviality defined above to define a family of IEC logics, each of which respects (some candidate for) triviality.

So, returning to the full signature with operators K and Δ, and with sphere models defined as in §2, I will now focus on specific model classes. This defines not one unique logic, but a family of related logics.

Definition 9 (IEC Model Classes) *For each $n\in\mathbb{N}$, let $\mathbb{C}_n$ be the class of all sphere models that respect the candidate set $\mathbb{T}_n$ as follows:*

$$\mathbb{C}_n=\{M:\textit{for all }\mathbb{T}_n\langle\Gamma,\varphi\rangle,(\Gamma\text{-}\varphi\text{-}\textit{RS})\textit{ holds at all }w\in W\textit{ of }M\}$$

For each candidate $\mathbb{T}_n$ above, we see that the IEC principles hold for *just those consequences* (considered as trivial) in the model class $\mathbb{C}_n$. The behavior of ‘determinate knowledge’, thus, depends purely on the structure of the notion of triviality. (cf. Jago 2014a, p.251, Theorem 8.3)

Corollary 1 *If $\mathbb{T}_n\langle\Gamma,\varphi\rangle$, then $\{\Delta K\psi:\psi\in\Gamma\}\vDash_{\mathbb{C}}\neg\Delta\neg K\varphi$.*

Proof. By an application of The Γ-ψ Respect Theorem. □

Since candidates for triviality increase in size $\mathbb{T}_n \subset \mathbb{T}_{n+1}$, as per the previous observation of Heredity, the corresponding model classes decrease in size $\mathbb{C}_{n+1} \subset \mathbb{C}_n$ and their consequence relations subsequently get stronger. In particular, the stronger logics in this family validate an increasing number of cases of the IEC principle. For example:

Remark 3 IEC principles for various $\mathbb{C}_n$-Consequence relations.

- $\Delta K(\varphi \to \psi), \Delta K\varphi \vDash_{\mathbb{C}_3} \neg\Delta\neg K\psi$ **(3-MP)**
- $\Delta K\neg\varphi \vDash_{\mathbb{C}_5} \neg\Delta\neg K(\varphi \to \bot)$ **(5-Neg)**
- $\Delta K(\varphi \to \psi), \Delta K(\psi \to \chi) \vDash_{\mathbb{C}_7} \neg\Delta\neg K(\varphi \to \chi)$ **(7-Trans)**
- $\Delta K(\varphi \to \psi) \vDash_{\mathbb{C}_9} \neg\Delta\neg K(\varphi \to (\psi \wedge \varphi))$ **(9-Pool)**

By way of illustration, consider the classically valid inference form of Transitivity and the 7-Trans result. According to these definitions, the narrowest candidates for triviality do not include Transitivity because it is not a paradigm case (not derivable by a single primitive rule). There are, however, more expansive candidate extensions of triviality that do consider the inference form Transivity to be trivial and for *any logic* that respects these candidates, the relevant IEC principle is valid (from $\mathbb{C}_7$ and above).

The family of $\mathbb{C}_n$-Consequence relations represent one fully worked example of the ranking approach to triviality. In this final section of the paper, I show how to axiomatize this family of logics.

5 Axiomatization

For each rank $n \in \mathbb{N}$, we can define a *logic of indeterminate closure* LIC$_n$ as an extension of the classical Hilbert system for CPL. These have a shared, core set of axioms and rules, but they each have different axioms licensed by (A8) as the relevant versions of the IEC principle.

Definition 10 (Axioms of LIC$_n$) *For each* $n \in \mathbb{N}$*, the set of* axioms *AX$_n$ is the set of all instances of the following schemata.*

(A1) $\varphi \to (\psi \to \varphi)$

(A2) $(\varphi \to (\psi \to \chi)) \to ((\varphi \to \psi) \to (\varphi \to \chi))$

(A3) $(\neg\varphi \to \neg\psi) \to (\psi \to \varphi)$

(A4) $K\varphi \to \varphi$

(A5) $\Delta\varphi \to \varphi$

(A6) $\Delta\varphi \to \neg\Delta\neg\Delta\varphi$

(A7) $\Delta(\varphi \to \psi) \to (\Delta\varphi \to \Delta\psi)$

(A8) $\Delta K\psi_1 \to (\ldots \to (\Delta K\psi_m \to \neg\Delta\neg K\varphi))$
for all reductive implications $\psi_1 \to (\ldots \to (\psi_m \to \varphi)) \in RI_n$[11]

Much as before, we use the axioms to define the primitive rules.

Definition 11 (Primitive Rules of LIC_n) *A* derivation from Γ *in* LIC_n *is a finite (possibly empty) sequence of sentences* $\psi_1, \ldots, \psi_n$. *N.B. in the rules below,* δ *is used as a variable for such sequences.*

(R1) *If* δ *is a derivation from* Γ *and* $\varphi \in \Gamma$, *then* δ, φ *is a derivation from* Γ.

(R2) *If* δ *is a derivation from* Γ *and* $\varphi \in AX_n$, *then* δ, φ *is a derivation from* Γ.

(R3) *If* δ *is a derivation from* Γ *and there are members of this sequence of the form* ψ *and* $\psi \to \varphi$, *then* δ, φ *is a derivation from* Γ.

(R4) *If* δ *is a derivation from* Γ *and there is a member of this sequence* φ *that does not depend on any* $\psi \in \Gamma$, *then* $\delta, \Delta\varphi$ *is a derivation from* Γ.

For the definition of derivability *per se*, we eliminate the empty sequence.

Definition 12 (Derivability in LIC_n) φ *is derivable from* Γ *in* LIC_n, *written* $\Gamma \vdash_{LIC_n} \varphi$, *iff there is a non-empty derivation from* Γ *ending in* φ, *i.e. a sequence of sentences* $\psi_1, \ldots, \psi_n$ *that satisfies the criteria of Def. 11 such that* $\psi_n = \varphi$. *This sequence* witnesses *that* $\Gamma \vdash_{LIC_n} \varphi$.

The completeness proof now follows by a canonical model construction. I will only sketch the interesting details. We first relativize the notions of a consistent and maximal set of sentences to each logic, LIC_n.

Definition 13 (n-Consistency) $Con_n(\Gamma)$ *iff for all sentences* φ *we have at least one of* $\Gamma \nvdash_{LIC_n} \varphi$ *or* $\Gamma \nvdash_{LIC_n} \neg\varphi$

[11] These are the sets RI_n from Definition 8, defined only over the signature of the base (classical) logic. Thanks to an anonymous referee for pointing out the potential ambiguity.

Definition 14 (n-Maximal-Consistency) *$MCS_n(\Gamma)$ iff $Con_n(\Gamma)$ and it is* not *the case that $Con_n(\Gamma \cup \{\varphi\})$ for any $\varphi \notin \Gamma$*

As usual, it follows that inconsistent sets can prove anything, and that $\Gamma \vdash_{\text{LIC}_n} \varphi$ iff it is *not* the case that $\text{Con}_n(\Gamma \cup \{\neg\varphi\})$. Maximal sets are deductively closed, and each logic supports Lindenbaum's Lemma.

Lemma 2 (Lindenbaum) *If $Con_n(\Gamma)$, then there is some sets of sentences Γ^* such that $\Gamma \subseteq \Gamma^*$ and $MCS_n(\Gamma^*)$.*

We define the proper canonical model and prove the Truth Lemma.

Definition 15 (The Proper Canonical Model of LIC_n) *Is a six element structure $\mathfrak{M}_n = \langle W, C, H, I, N, S\rangle$ with the following components.*

- $W_{\mathfrak{M}_n} = \{\Gamma : MCS_n(\Gamma)\}$
- $C_{\mathfrak{M}_n} = Sent$
- $H_{\mathfrak{M}_n}(\varphi) = \varphi$
- $I_{\mathfrak{M}_n}(\varphi) = \{\Gamma : \varphi \in \Gamma\}$
- $N_{\mathfrak{M}_n}(\Gamma) = \{\varphi : K\varphi \in \Gamma\}$
- $S_{\mathfrak{M}_n}(\Gamma) = \{\Sigma : \textit{for all } \varphi \in Sent, \textit{ if } \Delta\varphi \in \Gamma, \textit{ then } \varphi \in \Sigma\}$

Lemma 3 (Truth) $\Gamma \in [\![\varphi]\!]^{\mathfrak{M}_n}$ *iff* $\varphi \in \Gamma$

Proof. This is quick: $\Gamma \in [\![\varphi]\!]^{\mathfrak{M}_n}$ iff $\Gamma \in I_{\mathfrak{M}_n}(H_{\mathfrak{M}_n}(\varphi))$ by the definition of CGCs iff $\varphi \in \Gamma$ by the definitions of $H_{\mathfrak{M}_n}$ and $I_{\mathfrak{M}_n}$. □

However, it still remains to establish that this simplistic structure *really is* a sphere model and that it belongs to the intended model class. For the first part, we need to see that the composition of $H_{\mathfrak{M}_n}$ and $I_{\mathfrak{M}_n}$ is well-behaved, i.e. that is satisfies all desired operations on 'truth sets'.

Lemma 4 ($\mathfrak{M}_n$ is a Sphere Model)

Proof. I present the illustrative cases of $\neg, K, \Delta$.

- $\Gamma \in [\![\neg\varphi]\!]^{\mathfrak{M}_n}$
  iff $\neg\varphi \in \Gamma$ by the Truth Lemma
  iff $\varphi \notin \Gamma$ by consistency
  iff $\Gamma \notin [\![\varphi]\!]^{\mathfrak{M}_n}$ by the Truth Lemma

- $\Gamma \in [\![K\varphi]\!]^{\mathfrak{M}_n}$
  iff $K\varphi \in \Gamma$ by the Truth Lemma
  iff $\varphi \in N_{\mathfrak{M}_n}(\Gamma)$ by definition of $N_{\mathfrak{M}_n}$
  iff $H_{\mathfrak{M}_n}(\varphi) \in N_{\mathfrak{M}_n}(\Gamma)$ by definition of $H_{\mathfrak{M}_n}$

- $\Gamma \in [\![\Delta\varphi]\!]^{\mathfrak{M}_n}$
  iff $\Delta\varphi \in \Gamma$ by the Truth Lemma
  iff $\Sigma \in S_{\mathfrak{M}_n}(\Gamma)$ implies $\Sigma \in [\![\varphi]\!]^{\mathfrak{M}_n}$ by definition of $S_{\mathfrak{M}_n}$
  iff $S_{\mathfrak{M}_n}(\Gamma) \subseteq [\![\varphi]\!]^{\mathfrak{M}_n}$

□

Finally, we need to establish that this structure is in the right model class. That means: that the respect schema for all n-Trivial inference forms holds in all worlds (MCSs) of $\mathfrak{M}_n$. The following lemma will be useful.

Lemma 5 (Extension) *If MCS*$_n(\Gamma)$ *and* $\Delta\varphi \notin \Gamma$*, then in* $\mathfrak{M}_n$ *there is some* $\Sigma \in S_{\mathfrak{M}_n}(\Gamma)$ *such that* $\varphi \notin \Sigma$*. Contrapositively, we can infer that if all MCSs in the sphere of* Γ *in* $\mathfrak{M}_n$ *contain* φ*, then* $\Delta\varphi \in \Gamma$*.*

Proof. Let $\text{MCS}_n(\Gamma)$ and $\Delta\varphi \notin \Gamma$. I will first show that $\text{Con}_n(\Sigma')$ for the set $\Sigma' = \{\neg\varphi\} \cup \{\psi : \Delta\psi \in \Gamma\}$. Suppose not. Then there are finite $\psi_i \in \Sigma'$ such that $\vdash_{\text{LIC}_n} \psi_1 \to (\ldots \to (\psi_m \to \varphi))$. Since this sentence is derivable, it follows that $\vdash_{\text{LIC}_n} \Delta\psi_1 \to (\ldots \to (\Delta\psi_m \to \Delta\varphi))$ by the normality of Δ and so we have $\Gamma \vdash_{\text{LIC}_n} \Delta\varphi$. By hypothesis, however, we have $\Gamma \nvdash_{\text{LIC}_n} \Delta\varphi$, so indeed we have $\text{Con}_n(\Sigma')$ by reductio. We can apply Lindenbaum's Lemma to infer the existence of the target set: there is some Σ such that $\Sigma' \subseteq \Sigma$ and $\text{MCS}_n(\Sigma)$. By construction, if $\Delta\psi \in \Gamma$, then $\psi \in \Sigma$, so in the canonical model we have $\Sigma \in S_{\mathfrak{M}_n}(\Gamma)$ as desired, and $\neg\varphi \in \Sigma$. □

We show that the canonical model is in the right model class.

Lemma 6 (Rank) $\mathfrak{M}_n \in \mathbb{C}_n$

Proof. Let $\mathbb{T}_n\langle\Gamma, \varphi\rangle$. Note that there is at least one associated, reductive implication $\psi_1 \to (\ldots \to (\psi_m \to \varphi)) \in \text{RI}_n$ such that $\{\psi_1, \ldots, \psi_m\} \subseteq \Gamma$. Then by (A8) we have $\vdash_{\text{LIC}_n} \Delta K\psi_1 \to (\ldots \to (\Delta K\psi_m \to \neg\Delta\neg K\varphi))$, call this derivable sentence (*). Note that (*) is contained in all worlds of $\mathfrak{M}_n$. We can now show that (Γ-φ-RS) holds in all worlds.

Suppose that $H_{\mathfrak{M}_n}(\Gamma) \subseteq \bigcap\{N_{\mathfrak{M}_n}(\Sigma) : \Sigma \in S_{\mathfrak{M}_n}(\Gamma)\}$. Let $\Pi \in S_{\mathfrak{M}_n}(\Gamma)$. It is then easy to work out that $\{K\psi_1, \ldots, K\psi_m\} \subseteq \Pi$ (with respect to the sentences $\{\psi_1, \ldots, \psi_m\} \subseteq \Gamma$ from above). Since this holds throughout $S_{\mathfrak{M}_n}$, it follows by the Extension Lemma that $\{\Delta K\psi_1, \ldots, \Delta K\psi_m\} \subseteq \Gamma$. Then since (*) is contained in Γ we have $\neg\Delta\neg K\varphi \in \Gamma$ by deductive closure and hence $\Delta\neg K\varphi \notin \Gamma$. It follows by the Extension Lemma that there is some $\Sigma \in S_{\mathfrak{M}_n}(\Gamma)$ such that $\neg K\varphi \notin \Sigma$ and hence $K\varphi \in \Sigma$. Thus, we have $H(\varphi) \in N_{\mathfrak{M}_n}\Sigma$ and hence $H_{\mathfrak{M}_n}(\varphi) \in \bigcup\{N_{\mathfrak{M}_n}(\Sigma) : \Sigma \in S_{\mathfrak{M}_n}(\Gamma)$. □

It follows that the logic LIC_n is sound and complete.

Theorem 2 (Adequacy) $\Gamma \vDash_{\mathbb{C}_n} \varphi$ *iff* $\Gamma \vdash_{LIC_n} \varphi$

Proof. Soundness is straightforward. For completeness, we reason that if $\Gamma \nvdash_{\text{LIC}_n} \varphi$, then $\text{Con}_n(\Gamma \cup \{\neg\varphi\})$, so there is a world (MCS) of the canonical model $\Sigma \in W_{\mathfrak{M}_n}$ with $\Gamma \cup \{\neg\varphi\} \subseteq \Sigma$, thus by the Truth Lemma, we have a member of the model class $\mathbb{C}_n$ which shows that $\Gamma \nvDash_{\mathbb{C}_n} \varphi$. □

6 Conclusion

IEC logics offer a formal solution to an epistemic paradox, by describing how trivial consequences are borderline cases of knowledge closure. I developed a simple semantics for these logics, without appeal to substantive metaphysics, and showed that such logics are axiomatizable.

In the process of showing these results, however, a number of question may have been raised. As emphasized in §3, the hard question for IEC theory is how to understand the concept of triviality. I described one way that proof-theory may constrain this concept.

The ranking approach to triviality provides formal structure, in the form of candidate extensions, that can be used to axiomatize IEC logics, but this is entirely determined by the choice of *primitive rules* of the base logic. So, from a philosophical point of view, there is more to say about the actual primitive rules that we ‘really use’ or that ‘really give structure’ to our concept of triviality. This is a difficult and important question, but one that lies beyond the scope of the present paper.

References

Artemov, S., & Kuznets, R. (2014). Logical omniscience as infeasibility. *Annals of Pure and Applied Logic*, *165*(1), 6–25.

Cobreros, P., Egré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, *41*(2), 347–385.

Duc, H. N. (1995). Logical omniscience vs. logical ignorance: On a dilemma of epistemic logic. In C. Pinto-Ferreira & N. J. Mamede (Eds.), *Progress in Artificial Intelligence. EPIA 1995.* (pp. 237–248). Berlin, Heidelberg: Springer.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*(4), 265–300.

Hawke, P., Özgün, A., & Berto, F. (2020). The fundamental problem of logical omniscience. *Journal of Philosophical Logic*, *49*(4), 727–766.

Hintikka, J. (1962). *Knowledge and Belief: an Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.

Jago, M. (2006). Hintikka and Cresswell on logical omniscience. *Logic and Logical Philosophy*, *15*(4), 325–354.

Jago, M. (2014a). *The Impossible: An Essay on Hyperintensionality*. Oxford: Oxford University Press.

Jago, M. (2014b). The problem of rational knowledge. *Erkenntnis*, *79*(suppl. 6), 1151–1168.

Sedlár, I. (2021). Hyperintensional logics for everyone. *Synthese*, *198*(2), 933–956.

Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.

Colin R. Caret
Utrecht University, Department of Philosophy and Religious Studies
Netherlands
E-mail: colin.caret@protonmail.com