Problems in Epistemic Space*

Jens Christian Bjerring

forthcoming in *Journal of Philosophical Logic*

Abstract

When a proposition might be the case, for all an agent knows, we can say that the proposition is epistemically possible for the agent. In the standard possible worlds framework, we analyze modal claims using quantification over possible worlds. It is natural to expect that something similar can be done for modal claims involving epistemic possibility. The main aim of this paper is to investigate the prospects of constructing a space of worlds—epistemic space—that allows us to model what is epistemically possible for ordinary, non-ideally rational agents like you and me. I will argue that the prospects look dim for successfully constructing such a space. In turn, this will make a case for the claim that we cannot use the standard possible worlds framework to model what is epistemically possible for ordinary agents.

*I would especially like to thank Ben Blumson, Weng Hong Tang, and an anonymous referee from the *Journal of Philosophical Logic* for extensive comments on the paper. I would also like to thank David Chalmers, Alan Hájek, Hui Chieh Loy, Wolfgang Schwarz, Johanna Seibt, and audiences at Aarhus University and National University of Singapore for very helpful discussion of earlier versions of this paper.

1 Introduction

I wonder about a lot. I wonder whether it rains in Canberra, whether Shakespeare was really Edward de Vere, and whether Goldbach's Conjecture is false. As it happens, I lack the required evidence to determine whether these propositions are true. So for all I know, they might be true, and they might not.

When p might be true, for all an agent knows, we can say that p is epistemically possible for the agent. It is, amongst other things, epistemically possible for me that it rains in Canberra, that Shakespeare was really Edward de Vere, and that Goldbach's Conjecture is false. Similarly, when an agent knows p, we can say that p is epistemically necessary for the agent. It is, amongst other things, epistemically necessary for me that I am currently in Denmark, that Shakespeare is Shakespeare, and that $2 + 2 = 4$.

We are used to analyzing modal claims using quantification over possibilities of some sort, and it is natural to hope that something similar can be done for modal claims involving epistemic possibility. In particular, it is natural to hold that p is epistemically possible for an agent whenever there is a *scenario*—or an *epistemically possible world*—at which p is true. Corresponding to the epistemic possibility that Shakespeare really was Edward de Vere are then various scenarios at which Shakespeare originated in the gametes that Edward de Vere actually did—and various scenarios at which Shakespeare originated in the gametes that he actually did. And corresponding to the epistemic possibility that Goldbach's Conjecture is false are various scenarios at which there is some even integer greater than or equal to 4 that cannot be expressed as the sum of two primes. Call the space of all scenarios *epistemic space*.[1]

At the core of this familiar way of thinking about possibilities in a world-involving framework is the inverse relation between gaining information—covering both knowledge and belief—and excluding possibilities:[2]

> (**Information Principle**) For any gain in information, there is a corresponding elimination of scenarios—and vice versa.

[1]The intuitive picture of epistemic possibility presented above is influenced by Chalmers 2011; see also Hintikka 2003.

[2]Throughout the paper I will talk about knowledge and epistemic possibility rather than belief and doxastic possibility, but the main results apply mutatis mutandis to belief and doxastic possibility.

For instance, when I come to know that it rains in Canberra, all scenarios in epistemic space at which it does not rain in Canberra are eliminated as epistemically impossible for me. (Information Principle), I take it, is a constitutive feature of the Hintikka-inspired possible worlds framework for modeling knowledge, information, and belief.

While David Chalmers 2011, Jaakko Hintikka 1962, and David Lewis 1986—among several others—have advanced world-involving frameworks similar to the one described above to model what is epistemically possible for idealized reasoners, I want to argue that the prospects look dim for using the framework to model what is epistemically possible for ordinary reasoners like you and me. I will center my discussion around Mark Jago's model of epistemic space in his article "Logical Information and Epistemic Space" (cf. Jago 2009). Although Jago's model holds promise of giving us an account of epistemic possibility for ordinary agents, I will argue that it cannot do the required work. This criticism, in turn, will make a case for the claim that we must give up (Information Principle) to model epistemic possibility for ordinary agents, and hence that we cannot use the possible worlds framework to model epistemic possibility for such agents.

Here is the plan. In section 2, I lay out the role of epistemic space. In section 3, I present Jago's model of epistemic space. In section 4, I criticize Jago's model and use this criticism to support the claim that (Information Principle) must go to model what is epistemically possible for ordinary agents. In section 5, I conclude.

2 Epistemic Possibilities and Spaces

The general task is to construct an epistemic space that allows us to model claims of epistemic possibility—like the intuitive ones above—in terms of what is true at scenarios in this space. It is useful first to follow Chalmers and distinguish between *strict* and *deep* epistemic possibility:

> [T]he notion of *strict epistemic possibility*—ways things might be, for all we know—is undergirded by a notion of *deep epistemic possibility*—ways things might be, prior to what anyone knows. Unlike strict epistemic possibility, deep epistemic possibility does not depend on a particular state of knowledge, and is not obviously relative to a subject. Whereas it is strictly epistemically possible (for a subject) that p when there is some epistemically possible scenario (for that subject) in which p, it is deeply epistemically possible that p when there is some deeply epistemically possible scenario in which p. (Chalmers 2011, p. 62.)

Generally, the notion of deep epistemic possibility delineates the borders of epistemic space, and within it, strict epistemic possibility for an agent delineates the borders of what this agent knows. When an agent comes to know p, this piece of knowledge divides the class of deeply epistemically possible scenarios into those scenarios at which p is true and those at which p is false. If an agent knew nothing at all, all deeply epistemically possible scenarios would also be strictly epistemically possible for this agent.

More precisely, whereas the class of deeply epistemically possible scenarios constitute epistemic space W, the class of scenarios in W that remain strictly epistemically possible for an agent a (at a given time) constitute strict epistemic space W_a. The relationship between strict and deep epistemic possibility can then be captured by saying that for any agent a, $W_a \subseteq W$. So if w is in W_a, then w is in W, though the converse need not be the case. In this sense, deep epistemic possibility is a necessary, yet not a sufficient condition for strict epistemic possibility.

Since deep epistemic possibility is a necessary condition for strict epistemic possibility, our central task is hence to establish the following plentitude principle:[3]

> (**Plentitude**) For all sentences A, A is deeply epistemically possible iff there is a scenario w in W such that A is true at w.

If (Plentitude) is to do the required work, we must construct an epistemic space that contains enough scenarios to ensure that all sentences that remain strictly epistemically possible for ordinary agents are also deeply epistemically possible. For if we can guarantee the deep epistemic possibility of A, we will have guaranteed the potential strict epistemic possibility of A for some agent.

2.1 Deep Epistemic Possibility: Negative Characterization

How should we think of deep epistemic possibility? It depends. I am interested in developing an epistemic space that in part allows us to make sense of claims such as:

[3] The plentitude principle is taken from Chalmers 2011. For the purposes of this paper, I take the objects of strict and deep epistemic possibility to be *sentences*. Since issues about context-dependence will not play a role in my arguments, I will take these sentences to be declarative sentence *types* in a language $\mathcal{L}$ that has no context-dependent expressions. $\mathcal{L}$, however, has symbols that play the same inferential roles as the logical connectives in standard propositional logic. If one has no quarrels with primitive, arbitrarily fine-grained propositions, one may translate everything I say about sentences into talk about such propositions.

(**Shakes**) It is strictly epistemically possible for agent a that Shakespeare was really Edward de Vere.

(**Math**) It is strictly epistemically possible for agent a that Goldbach's Conjecture is false.

(**Logic**) It is strictly epistemically possible for agent a that (some tautology) $((A \leftrightarrow (B \wedge \neg A)) \rightarrow \neg B)$ is false.

To develop a model for these kinds of strict epistemic possibility claims, it is easy to see what deep epistemic possibility *cannot* be.

On the assumption that Kripke is right about the necessity of origin, deep epistemic possibility cannot be *metaphysical possibility*. For, on this assumption, "Shapespeare is Edward de Vere" cannot metaphysically possibly be true. But although metaphysically impossible, "Shakespeare is Edward de Vere" is nevertheless strictly epistemically possible for some ordinary agent. Since deep epistemic possibility is a necessary condition for strict epistemic possibility, deep epistemic possibility cannot be metaphysical possibility.[4] Something similar holds for (Math) and (Logic). Since all a priori falsehoods are metaphysically impossible, we cannot identify deep epistemic possibility with metaphysical possibility to model (Math) and (Logic).

Deep epistemic possibility cannot be *idealized a priori possibility*. I have in mind Chalmers' conception of idealized apriority:

> This idealized notion of apriority abstracts away from contingent cognitive limitations. If there is any possible mental life that starts from a thought and leads to an a priori justified acceptance of that thought, the thought is a priori. So if a hypothesis can be known to be false only by a great amount of a priori reasoning, it is nevertheless deeply epistemically impossible. [...] As a result, this idealization is best suited for modeling the knowledge and belief of idealized reasoners that may be empirically ignorant, but that can engage in arbitrary a priori reasoning. (Chalmers 2011, p. 66.)

On this picture, neither Goldbach's Conjecture—assuming that it is indeed true—nor $((A \leftrightarrow (B \wedge \neg A)) \rightarrow \neg B)$ can ideally a priori possibly be false. But since the falsity of these sentences may be strictly epistemically possible for some ordinary agents, deep epistemic possibility cannot be idealized a priori possibility. Yet, we can use idealized

[4]Robert Stalnaker 1984, however, has argued that we can understand the belief and knowledge of ordinary agents by appeal only to metaphysical possibilities. I will not evaluate Stalnaker's proposal in this paper, but see Robbins 2004 for discussion.

a priori possibility to make sense of (Shakes). For on an intuitive view of apriority, "Shakespeare is Edward de Vere" is not a priori false.

Deep epistemic possibility cannot be *logical possibility.* For $\neg((A \leftrightarrow (B \wedge \neg A)) \rightarrow \neg B)$ cannot logically possibly be true. But since such logical falsities may be strictly epistemically possible for some ordinary agents, deep epistemic possibility cannot be logical possibility. Yet, we can use logical possibility to make sense of (Shakes) and (Math). For at least on a narrow understanding of logical possibility, there is nothing in logic per se that rules out the falsity of Goldbach's Conjecture.

Since (Shakes), (Math), and (Logic) describe strict epistemic possibilities for some ordinary agents, deep epistemic possibility cannot be identified with metaphysical, idealized a priori, or logical possibility. Rather, the relevant notion of deep epistemic possibility has to be more permissive than these notions of possibility. In particular, some deep epistemic possibilities must be *logically impossible*—and hence metaphysically and idealized a priori impossible.

2.2 Deep Epistemic Possibility: Positive Characterization

To define a useful notion of deep epistemic possibility, let me first be more precise about the kinds of agents that I am interested in. In a broad sense, I am interested in what may remain strictly epistemically possible for ordinary agents who are capable of rational—although not ideally rational—reasoning. Ordinary agents, I take it, have a basic capacity for non-trivial inferential reasoning. Although they often fail to tease out *all* the consequences of what they know, they often tease out at least the *obvious* consequences of what they know.[5] And although they occasionally fail to tease out even the obvious consequences of their knowledge, it remains true that they have the dispositional capacity to do so. For instance, even if I get distracted by a nearby passing car and do not infer that I will get wet from my knowledge that it just started to rain and that if it rains then I will get wet, my dispositional capacity to infer the former piece of information from the latter remains in place.

In a narrower sense, I will be focusing on what may remain strictly epistemically possible for *minimally rational agents* who have limited, but non-trivial inferential capacities for (a priori) *propositional logical* reasoning. Such agents are *rational* because they always tease out the obvious logical consequences of what they know, and they

[5]For motivation of similar constraints on minimal rationality, see Cherniak 1986, Harman 1999, and Field 2001.

are *minimally* rational because they only have limited or bounded capacities available for (propositional) logical reasoning. In particular, minimally rational agents always know that obvious logical truths such as $(A \to A)$ are true and that obvious logical falsities such as $(A \wedge \neg A)$ are false. And they know these facts—like most ordinary agents do on reflection—because it only requires basic logical acumen to realize that $(A \to A)$ is tautological and that $(A \wedge \neg A)$ is inconsistent.

I take logical reasoning to be governed by *rules*.[6] In particular, minimally rational agents engage in logical reasoning by applying various inference rules governing the propositional connectives such as $\neg$ and $\to$, and they engage in complex chains of logical reasoning by applying these rules successively. Let $\mathcal{R}$ be a complete set of inference rules in propositional logic that minimally rational agents can apply. For reasons that will become clear in sections 4.2 and 4.3, agents must in general be able to reason logically from the empty set of premises. So effectively, some rules in $\mathcal{R}$ will count as axioms (schemas).[7] But given this, my results are compatible with different choices for the relevant system of inference rules. One might, for instance, adopt one of the axiomatic proof systems for propositional logic discussed in Mendelson 1997 (pp. 33-49), or a propositional sequent calculus akin to the one discussed in Buss 1998 (pp. 10-18).

To make the idea of limited or bounded logical reasoning precise, I invoke the idea of *step-based* logical reasoning.[8] Consider, for instance, the inference from $(A \wedge B)$ to A and the inference from $((A \wedge B) \wedge C)$ to A. Whereas the former takes 1 step of logical reasoning using conjunction elimination once, the latter takes 2 steps of logical reasoning using conjunction elimination twice. Limited or bounded logical reasoning can then be characterized in terms of *n-step logical reasoning*, where n is a finite natural number, and where 1 step of logical reasoning corresponds to one application of an inference rule in $\mathcal{R}$ to a set of sentences. Intuitively, I intend n to be a measure of the amount of cognitive resources that an agent has available for logical reasoning—or alternatively, a measure of an agent's logical competence.

We can then say that minimally rational agents are able to perform up to n steps

[6]For views that take rational logical reasoning, broadly construed, to be governed by rules or inferential roles, see Field 1977 and Wedgwood 2007.

[7]I will not consider systems of inference rules that allow agents to reason logically from assumptions. As far as I can tell, though, there is nothing that should prevent my general results from applying to agents who reason in such systems as well.

[8]For similar ideas, see Drapkin and Perlis 1986.

of logical reasoning, where n corresponds to the number of applications of rules in $\mathcal{R}$ required to disprove obvious logical falsities such as $(A \wedge \neg A)$ and to prove obvious logical truths such as $(A \rightarrow A)$. This captures the idea that minimally rational agents always know that obvious logical falsities are false—are able to disprove such falsities—and that obvious logical truths are true—are able to prove such truths. While it is unclear which logical truths are obvious, it is clear that all truths that can be established by 1 step of logical reasoning are obvious. So if a precise value is wanted for the number of logical steps that a minimally rational agent can perform, we can take it to be 1. In turn, we can take an obvious logical truth (falsity) to be any truth (falsity) that can be proved (disproved) in 1 step of logical reasoning.

Given this, we can now spell out deep epistemic necessity and possibility by appeal to *provability in n steps of logical reasoning using the rules in* $\mathcal{R}$. To that end, let a proof of A in n steps of logical reasoning be a derivation of A from a set Γ of sentences—potentially the empty set—consisting of at most n applications of the rules in $\mathcal{R}$. Let a disproof of A in n steps of logical reasoning be a derivation of $\neg A$ from A—or from the set Γ of sentences such that $A \in \Gamma$—consisting of at most n applications of the rules in $\mathcal{R}$. Similarly, let a set Γ of sentences be disprovable in n steps of logical reasoning whenever there is a derivation of A and $\neg A$ from Γ consisting of at most n applications of the rules in $\mathcal{R}$. For simplicity, I will assume that agents can rule out sets of sentence that contain $\{A, \neg A\}$ non-inferentially. Finally, let '$\Box_n$' and '$\Diamond_n$' be metalinguistic operators, where '$\Diamond_n$' is defined as $\neg\Box_n\neg$. Read '$\Box_n A$' as 'A is provable in n steps of logical reasoning using the rules in $\mathcal{R}$', and read '$\Diamond_n A$' as 'A is not disprovable in n steps of logical reasoning using the rules in $\mathcal{R}$'. We can then define:

(**Deep-Nec**$_n$) A sentence A is deeply$_n$ epistemically necessary iff $\Box_n A$.

(**Deep-Pos**$_n$) A sentence A is deeply$_n$ epistemically possible iff $\Diamond_n A$.

Since n in (Deep-Nec$_n$) and (Deep-Pos$_n$) can take on a spectrum of values, we hence get a spectrum of notions of deep epistemic necessity and possibility.[9] For any finite n, there will be certain logical falsities that cannot be disproved within n steps of logical reasoning, and there will be certain logical truths that cannot be proved within n steps of logical reasoning. In turn, we can ensure that some deep epistemic

[9] Cf. Chalmers 2011, pp. 103-104.

possibilities will be logically impossible: for any finite n, some but not all logical falsities will be deeply$_n$ epistemically impossible, and some but not all logical truths will be deeply$_n$ epistemically necessary.

The n-step definition of deep epistemic possibility is tailor-made to the task at hand: to find a construction of an epistemic space that allows us to model strict epistemic possibility for minimally rational agents who can perform up to n steps of logical reasoning, for some sufficiently small n. For if we can establish (Deep-Nec$_n$) and (Deep-Pos$_n$), we will have ensured two things. First, we will have ensured that no obvious logical falsities will remain strictly epistemically possible for minimally rational agents. For given that obvious logical falsities can be disproved in n steps of logical reasoning, for some sufficiently small n, they will be deeply$_n$ epistemically impossible. Intuitively, this would capture the idea that minimally rational agents always know that obvious logical falsities are false. Second, we will have ensured that all obvious logical truths will remain strictly epistemically necessary for minimally rational agents. For given that obvious logical truths can be proved in n steps of logical reasoning, for some sufficiently small n, they will be deeply$_n$ epistemically necessary. Intuitively, this would capture the idea that minimally rational agents always know that obvious logical truths are true.

To vindicate a principle like (Plentitude) in the current setting, we need a spectrum of epistemic spaces $W_1, W_2, \ldots, W_n$ that can count as models for the spectrum of notions of deep epistemic possibility $\Diamond_1, \Diamond_2, \ldots, \Diamond_n$. In particular, we must look to establish:

> (**Plentitude**$_n$) For all sentences A, $\Diamond_n A$ iff A is true at some scenario w in W_n.
>
> (**N-Plentitude**$_n$) For all sentences A, $\Box_n A$ iff A is true at all scenarios w in W_n.

Jago's model of epistemic space—to which I turn now—holds promise of providing a space of scenarios that can make (Plentitude$_n$) and (N-Plentitide$_n$) true.

3 Jago's Epistemic Space

Jago's model $\mathcal{J}$ of epistemic space is a tuple:

$$\langle W^C, W^O, \nu, \preceq \rangle$$

W^C is the class of logically *closed* scenarios "at which the truths are both deductively closed and consistent" in the sense of classical logic. (Jago 2009, p. 334.) We can think of W^C as the class of logically possible scenarios. W^O is the class of *open* scenarios at which the truths "are inconsistent (and hence not closed)." (Jago 2009, p. 334.) Open scenarios were introduced by Graham Priest 2005 as a species of logically impossible worlds.[10] For *any* logically true sentence A, there are open scenarios at which A is false. Indeed, for *any* set of sentences, however logically inconsistent, there are open scenarios at which the truths include every member of this set. For my initial purposes, I assume—seemingly in line with Jago—that all open scenarios in W^O are logically inconsistent. Since any inconsistent set of sentences that is closed under (classical) logical deduction is extensionally equivalent to the set of all sentences, what is true according to scenarios in W^O should not be closed under logical deduction. For we need to be able to distinguish between what is true at different open scenarios.[11] ν is "a propositional valuation function, assigning a truth-value to each sentence at each [scenario]." (Jago 2009, p. 335.) Finally, $\preceq$ is a total order on $W^C \cup W^O$.

According to Jago, $w \preceq w'$ holds whenever the following holds: "if we expect agents to reject (some truth according to) w' a priori, then we also expect agents to reject (some truth according to) w a priori." (Jago 2009, p. 334.) While the maximal elements with respect to $\preceq$ are the logically possible scenarios in W^C, the minimal elements with respect to $\preceq$ are those scenarios in W^O at which "some obvious a priori impossibility is true, where that impossibility is as basic as an a priori impossibility can be." (Jago 2009, p. 334.) Let $\mathcal{X}$ denote a set of such obvious a priori impossibilities, and let $|w|$ denote the set of truths according to a scenario w. When $A \in |w|$, A is *true* at w, and I will also say that w *verifies* A. When $\neg A \in |w|$, A is *false* at

[10] For more on the nature of impossible worlds, see Bjerring forthcoming, Nolan 1997, and Priest 2005.

[11] Notice that if each w in W^O is maximal—in the sense of (Maximality) in section 4—then w is logically inconsistent if w is not (if the truths at w are not) deductively closed. But if some scenarios in W^O are not maximal—in the sense of either section 4.2 or section 4.3—it is possible to hold that all scenarios in W^O fail to be deductively closed but that only some such scenarios are logically inconsistent. Since it is not quite clear from Jago's presentation whether *all* scenarios in epistemic space are maximal or not, it is not quite clear whether all scenarios in W^O are logically inconsistent or merely not deductively closed. Yet, since my central arguments will have bite against either conception of what open scenarios in W^O are, the finer details need not worry us here.

w, and I will also say that w *falsifies* A.[12] We then have: if $\mathcal{X} \cap |w| \neq \emptyset$, $w \preceq w'$, for any scenario w'. That is, if we expect an agent to reject anything at all a priori, we expect her to reject all scenarios that verify some obvious a priori impossibility. I will assume that explicit contradictions of the form $\{A, \neg A\}$ are in $\mathcal{X}$. Given my focus on logical reasoning, we can then think of the minimal elements with respect to $\preceq$ as those logically inconsistent scenarios at which an explicit contradiction is true.

To determine when $\preceq$ holds in general, Jago invokes the idea of n-step reasoning. Model theoretically, he interprets the set of inference rules $\mathcal{R}$ as a binary relation $[[\mathcal{R}]]$ between scenarios.[13] In general, $(w, w') \in [[\mathcal{R}]]$ if and only if there is an instance of a rule in $\mathcal{R}$

$$\frac{A_1, A_2, \ldots, A_k}{B}$$

such that $\{A_1, A_2, \ldots A_k\} \subseteq |w|, B \notin |w|$, and $|w'| = |w| \cup \{B\}$. That is, when $(w, w') \in [[\mathcal{R}]]$, there is an instance of a rule schema in $\mathcal{R}$ that takes us from w to w', where w' verifies everything w does plus the conclusion B obtained by applying the rule once.

Jago then specifies a function f from open scenarios in W^O to $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$, where $f(w)$ gives the smallest (finite) number n of inference steps required to derive a basic a priori impossibility from $|w|$. For any two open scenarios w and w' such that $f(w) = m$ and $f(w') = n$, for $m < n$, we have that $w \preceq w'$. In words, if we expect an agent to rule out w' in n steps of logical reasoning, then we also expect the agent to rule out any scenario w from whose truths a basic a priori impossibility can be proved in less than n steps of logically reasoning.

This summarizes Jago's model. *If* Jago's model works, we can use it to construct a sequence of epistemic spaces as follows:

$$W_0 = \{w \in W^C \cup W^O | \text{ not: } f(w) = 0\}$$

[12] Although Jago does not specify what it means for a sentence to be false at a scenario w, it is natural to spell out falsity of A at w in terms of truth of $\neg A$ at w. If all scenarios are maximal—in the sense of (Maximality) in section 4—notice that A is false at w when $A \notin |w|$.

[13] Concerning the rules in $\mathcal{R}$, Jago says that they correspond to basic inferences that we expect rational agents to perform "in accordance with the meanings of 'and', 'or', 'for all' and so on." (Jago 2009, p. 335.) For my purposes, as mentioned, I focus solely on rules that govern the propositional connectives.

$$W_1 = \{w \in W^C \cup W^O | \text{ not: } f(w) \leq 1\}$$

$$W_2 = \{w \in W^C \cup W^O | \text{ not: } f(w) \leq 2\}$$

$$\vdots$$

$$W_n = \{w \in W^C \cup W^O | \text{ not: } f(w) \leq n\}$$

In words, for any finite n, W_n is the class of logically consistent and inconsistent scenarios from which a basic impossibility cannot be proved in up to n steps of logical reasoning. Given that explicit contradictions are in $\mathcal{X}$, there will be no scenarios in W_n that verify sentences or sets of sentences from which $\{A, \neg A\}$ can be proved within n steps of logical reasoning using the rules in $\mathcal{R}$. That is, there will be no scenarios in W_n that verify sentences—or sets of sentences—that can be disproved in up to n steps of logical reasoning.

If Jago's model works, we can use the spectrum of epistemic spaces to model (Plentitude$_n$): $\Diamond_n A$ iff A is true at some scenario w in W_n. To give a rough illustration of how, assume first $\Diamond_n A$. Then A is not disprovable in n steps. Suppose, for reductio, that $A \notin |w|$, for all scenarios $w \in W_n$. By construction of W_n, we can then infer a contradiction from any $|w|$ such that $A \in |w|$. But since A is not disprovable in n steps, there are plenty of scenarios in W_n whose truth-sets include A and from which a contradiction cannot be derived in n steps. So when $\Diamond_n A$, A is true at some scenario in W_n. Assume next that A is true at some $w \in W_n$. By construction of W_n, we cannot infer a contradiction from what is true according to w in n steps. So w does not verify any sentences that are disprovable within n steps of logical reasoning. So A is not disprovable in n steps of logical reasoning. So when A is true at some scenario in W_n, $\Diamond_n A$.

Derivatively, *if* Jago's model works, we can use it to construct a non-trivial epistemic space that contains enough scenarios to capture strict epistemic possibility for minimally rational agents. So the question remains: does Jago's model of epistemic space work? I do not believe that it does, and I shall argue why in the next section.

4 Problems in Jago's Epistemic Space

Generally, it is not quite clear from Jago's presentation whether scenarios in his model are maximal or not:

(**Maximality**) For all sentences A and scenarios $w \in W^C \cup W^O$, either $A \in |w|$ or $\neg A \in |w|$.

When a scenario obeys (Maximality), I will say that the scenario is *maximal*. If a scenario is not maximal, I will say that it is *non-maximal* or *partial*. When evaluated at partial scenarios, certain sentences are neither true nor false.

If a contradiction can be proved within n steps of logical reasoning from what is true according to a scenario, we know that a minimally rational agent can rule out—or reject a priori—the scenario in question. But if we allow partial scenarios in epistemic space, there is another type of logically "defective" scenarios that we might want minimally rational agents to be able to rule out as well. Consider, for instance, a partial scenario $w^\star$ at which only $(A \wedge B)$ is true but all other sentences neither true nor false. Although no contradiction can be proved from $|w^\star|$ using the rules in $\mathcal{R}$, there still seems to be a sense in which $w^\star$ is logically defective. For $w^\star$ fails to verify certain sentences that follow from $|w^\star|$ by basic logical reasoning—notably A and B that follow from $|w^\star|$ by two simple steps involving conjunction elimination. Intuitively, a minimally rational agent can easily realize that the description of the world that $w^\star$ offers is logically incomplete: it fails to describe something as true that can be established *is* true by basic logical reasoning. Call scenarios such as $w^\star$ *defective partial scenarios*—I will give a precise definition of these scenarios in section 4.2.

To fill in the details of Jago's model, we need first decide whether all scenarios in epistemic space are maximal. If we decide to include partial scenarios in epistemic space, we need then decide whether defective partial scenarios—in addition to scenarios from which a contradiction is provable—can be ruled out using rules in $\mathcal{R}$. Accordingly, there seems to be three ways of thinking about $\mathcal{J}$:

($\mathbf{C}_1$) All scenarios in $\mathcal{J}$ are maximal.

($\mathbf{C}_2$) Scenarios in $\mathcal{J}$ can fail to be maximal but some defective partial scenarios can be ruled out using rules in $\mathcal{R}$.

($\mathbf{C}_3$) Scenarios in $\mathcal{J}$ can fail to be maximal but no defective partial scenarios can be ruled out using rules in $\mathcal{R}$.

Corresponding to each of these possibilities, I will give three arguments. If $\mathcal{J}$ satisfies (C_1), the first argument shows that a contradiction can be proved from any logically inconsistent scenario in just 1 step of logical reasoning using the rules in $\mathcal{R}$. If $\mathcal{J}$ satisfies (C_2), the second argument shows that any scenario that fails to be logically consistent can be ruled out in just 1 step of logical reasoning using the rules in $\mathcal{R}$. These two arguments will show that $\mathcal{J}$ cannot be used to establish (Plentitude$_n$), but also that $\mathcal{J}$ wrongly characterizes minimally rational agents as logically omniscient. If $\mathcal{J}$ satisfies (C_3), the third argument shows that $\mathcal{J}$ cannot serve as an epistemic space for minimally rational agents. The third argument, in turn, will make a case for the claim that we must give up (Information Principle) to model strict epistemic possibility for minimally rational agents. I go through each argument in turn.

4.1 Argument against (C_1)

For the first argument, let us assume that scenarios in $\mathcal{J}$ satisfies (C_1) and hence that all scenarios in $W^C \cup W^O$ are maximal. Then $\neg A \in |w|$ whenever $A \notin |w|$. We know that w is logically inconsistent when a contradiction in $\mathcal{X}$ can be proved from $|w|$ using the rules in $\mathcal{R}$.

Given this, I will now prove (Result 1), where '$|w| \vdash_{\mathcal{R}} B$' abbreviates '$B$ is provable from the truths according to w using the rules in $\mathcal{R}$':

> (**Result 1**) For all sentences B and all scenarios w for which (C_1) holds, if $B \notin |w|$ and $|w| \vdash_{\mathcal{R}} B$, then $f(w) = 1$.

(Result 1) states that a contradiction can be proved from any logically inconsistent scenario in just 1 step using the rules in $\mathcal{R}$.

Let $[[\mathcal{X}]]$ be the class of open scenarios such that $\mathcal{X} \cap |w| \neq \emptyset$, and let '$|w| \vdash^n_{\mathcal{R}} B$' abbreviate '$B$ is provable from the truths according to w in n steps of logical reasoning using the rules in $\mathcal{R}$'. I prove (Result 1) by induction on the shortest number of steps required to prove a contradiction from $|w|$:[14]

> **Base case**. Assume $B \notin |w|$ and $|w| \vdash^1_{\mathcal{R}} B$. Since $|w| \vdash^1_{\mathcal{R}} B$, there is a w_1 such that $(w, w_1) \in [[\mathcal{R}]]$, where $|w_1| = |w| \cup B$. So $B \in |w_1|$. But also, since $B \notin |w|$, $\neg B \in |w|$ by (Maximality), and hence $\neg B \in |w_1|$.

[14]For the proof of (Result 1), I take the trivial inference rule that permits inferring A from any set Γ such that $A \in \Gamma$ as a 0-step inference rule. If there are independent reasons for not allowing such 0-step inference rules—perhaps because all reasoning, however basic and trivial, requires the use of some cognitive resources—(Result 1) will hold for $f(w) = 2$ rather than $f(w) = 1$.

So $w_1 \in [[\mathcal{X}]]$, and B and $\neg B$ can be proved from $|w|$ in 1 step of logical reasoning. So $f(w) = 1$ and (Result 1) holds for the base case.

Inductive step. Assume for the induction hypothesis that (Result 1) holds for all sentences B such that $|w| \vdash^n_{\mathcal{R}} B$. We want to show that (Result 1) holds when $|w| \vdash^{n+1}_{\mathcal{R}} B$. So let $n > 0$, and assume $B \notin |w|$ and $|w| \vdash^{n+1}_{\mathcal{R}} B$. Then B has to be proved from certain assumptions $A_1, A_2, \ldots A_k$ such that each A_i is provable from $|w|$ in at most n steps and such that B is provable from $A_1, A_2, \ldots A_k$ in 1 step. To show that $f(w) = 1$, there are two cases:

Case 1: For some A_i, $A_i \notin |w|$. By the induction hypothesis, if $A_i \notin |w|$ and $|w| \vdash^n_{\mathcal{R}} A_i$, then $f(w) = 1$.

Case 2: For all A_i, $A_i \in |w|$. We now repeat the argument from the base case. Since B can be derived from $A_1, A_2, \ldots, A_k$ in 1 step, and since each $A_i \in |w|$, $|w| \vdash^1_{\mathcal{R}} B$. Then there is a w_1 such that $(w, w_1) \in [[\mathcal{R}]]$, where $|w_1| = |w| \cup B$. So $B \in |w_1|$. But also, since $B \notin |w|$, $\neg B \in |w|$ by (Maximality), and hence $\neg B \in |w_1|$. So $w_1 \in [[\mathcal{X}]]$, and B and $\neg B$ can be proved from $|w|$ in 1 step of logical reasoning. So $f(w) = 1$.

So (Result 1) holds for the inductive step, and I conclude that (Result 1) holds in general. ∎

(Result 1) says that a contradiction can be proved from any logically inconsistent scenario in just 1 step of logical reasoning using the rules in $\mathcal{R}$. Effectively, (Result 1) shows that the attempt to construct a stratified epistemic space based on Jago's model will collapse when $\mathcal{J}$ satisfies (C_1). Since $f(w) = 1$, for all logically inconsistent scenarios in W^O, all scenarios that are not in W_0 will be logically consistent.

(Result 1) implies that we cannot use $\mathcal{J}$ to establish (Plentitude$_n$): $\Diamond_n A$ iff A is true at some scenario w in W_n. For suppose $\Diamond_m A$, for some A that is disprovable in n steps, for some n such that $n > m > 1$. That is, assume that A—although disprovable in n steps of logical reasoning—is not disprovable in m steps, for some m greater than 1 but smaller than n. To establish (Plentitude$_n$), we need a scenario $w \in W_m$ at which A is true. By (Result 1), however, there are no logically inconsistent scenarios in any W_n, for $n \geq 1$. So A is not true at any $w \in W_m$. So $\Diamond_m A$ is true but there is no scenario in W_m at which A is true. Thus (Plentitude$_n$) fails.

Furthermore, (Result 1) implies (Omni$_1$):

> (**Omni$_1$**) Any agent who can perform just 1 step of logical reasoning using the rules in $\mathcal{R}$ can a priori reject any logically inconsistent scenario in $W^C \cup W^O$.

Given (Omni$_1$), it follows that any minimally rational agent capable of performing just 1 step of logical reasoning can rule out any logically inconsistent scenario. As a result, only logically consistent scenarios remain strictly epistemically possible for such agents, and derivatively, all logical truths are strictly epistemically necessary for them. If we follow the Hintikkian tradition and take knowledge to be a kind of (strict) epistemic necessity, all minimally rational agents are then wrongly characterized as logically omniscient.

This concludes my first argument. To avoid (Result 1), scenarios in $\mathcal{J}$ cannot be maximal.

4.2 Argument against (C$_2$)

For the second argument, let us assume instead that $\mathcal{J}$ satisfies (C$_2$). Then scenarios in $\mathcal{J}$ can fail to be maximal and some defective partial scenarios can be ruled out using rules in $\mathcal{R}$. When scenarios can fail to be maximal, we need *indeterminacies* in our framework. So let us say that:

> A is *true* at w just in case $A \in |w|$ (denoted by '$1(A)$ at w').
>
> A is *false* at w just in case $\neg A \in |w|$ (denoted by '$0(A)$ at w').
>
> A is *indeterminate* at w just in case neither $A \in |w|$ nor $\neg A \in |w|$ (denoted by '$\frac{1}{2}(A)$ at w').

If a contradiction can be proved from $|w|$ within n steps of logical reasoning using the rules in $\mathcal{R}$, then a minimally rational agent can rule out the scenario w—whether w is maximal or partial. But to rule out defective partial scenarios, as discussed above, we need a slightly more elaborate account. We can say that a partial scenario w is *defective* just in case A is indeterminate at w but either A or $\neg A$ can be proved from $|w|$ using the rules in $\mathcal{R}$. In turn, we can say that a minimally rational agent can rule out a scenario w either when A is indeterminate at w but A or $\neg A$ can be proved from $|w|$ within n steps of logical reasoning, or when a contradiction can be proved from $|w|$ within n steps of logical reasoning.

To have a notationally convenient way of representing what it means for a minimally rational agent to rule out a scenario, I will use elements from the metalanguage. In particular, I will stipulate that $\{1(A), \frac{1}{2}(A)\}$ and $\{0(A), \frac{1}{2}(A)\}$ in addition to $\{1(A), 0(A)\}$ count as contradictions. Call this enriched notion of a contradiction an *i-contradiction*. We can then say that a partial scenario w is defective when $\frac{1}{2}(A)$ at w but either $1(A)$ or $0(A)$ can be proved from $|w|$ using $\mathcal{R}$—that is, when either the i-contradictions $\{\frac{1}{2}(A), 1(A)\}$ or $\{\frac{1}{2}(A), 0(A)\}$ can be proved from $|w|$ using $\mathcal{R}$. Notice that this does not mean that we have special inference rules that operate on the expressions '$1(A)$', '$0(A)$', and '$\frac{1}{2}(A)$' in the metalanguage. Inference rules operate on sentences in the object language, and the current notation is merely introduced for convenience.

In turn, we can say that a minimally rational agent can rule out a scenario w either when $\frac{1}{2}(A)$ at w but $1(A)$ or $0(A)$ can be proved from $|w|$ within n steps of logical reasoning, or when $\{1(A), 0(A)\}$ can be proved from $|w|$ within n steps of logical reasoning. In short, a minimally rational agent can rule out a scenario w just in case an i-contradiction can be proved from $|w|$ within n steps of logical reasoning using the rules in $\mathcal{R}$. To illustrate, suppose $\{A, (A \rightarrow B)\} \subseteq |w|$ but that both $B \notin |w|$ and $\neg B \notin |w|$. Then w is non-maximal and $\frac{1}{2}(B)$ at w. Since B can be proved from $|w|$ by just 1 application of modus ponens, an i-contradiction $\{\frac{1}{2}(B), 1(B)\}$ can be derived from $|w|$ in just 1 step. Any minimally rational agent who can perform 1 step of logical reasoning using $\mathcal{R}$ will then able to rule out w.

Finally, we can say that a scenario w *fails to be logically consistent* just in case an i-contradiction can be proved (simpliciter) from $|w|$ using the rules in $\mathcal{R}$. If w is maximal, w fails to be logically consistent whenever w is logically inconsistent—whenever $\{1(A), 0(A)\}$ can be proved from $|w|$ using $\mathcal{R}$. If w is non-maximal, w fails to be logically consistent just in case either $\{1(A), \frac{1}{2}(A)\}$, $\{0(A), \frac{1}{2}(A)\}$, or $\{1(A), 0(A)\}$ can be proved (simpliciter) from $|w|$ using $\mathcal{R}$.

Given this picture, $\mathcal{J}$ can satisfy (C_2). But if it does, we can also immediately prove a result similar to (Result 1). Omitting unnecessary formalism, the result is

(Result 2):[15]

> (**Result 2**) For all sentences B and all scenarios w for which (C_2) holds, if $0(B)$ at w or $\frac{1}{2}(B)$ at w and $|w| \vdash_{\mathcal{R}} B$, then an i-contradiction can be proved from $|w|$ in 1 step using the rules in $\mathcal{R}$.

(Result 2) states that an i-contradiction can be proved from any scenario that fails to be logically consistent in just 1 step using the rules in $\mathcal{R}$.

Since the proof strategy for (Result 2) is identical to the one for (Result 1), I will only show a part of the proof of (Result 2)—with a translation of symbolism, (Result 1) can also be seen as a special case of (Result 2) where we ignore the case concerning indeterminacies. I prove part of (Result 2) by induction on the shortest number of steps required to prove an i-contradiction from $|w|$:

> **Base case**. Assume $\frac{1}{2}(B)$ at w and $|w| \vdash^1_{\mathcal{R}} B$. Then B is provable from $|w|$ in 1 step. Since neither $B \in |w|$ nor $\neg B \in |w|$, $\frac{1}{2}(B)$ at w. So an i-contradiction ($\{1(B), \frac{1}{2}(B)\}$) can be proved from $|w|$ in 1 step using $\mathcal{R}$. Omitting the analogous reasoning for $0(B)$ at w, (Result 2) holds for the base case.
>
> **Inductive step**. Assume for the induction hypothesis that (Result 2) holds for all sentences B such that $|w| \vdash^n_{\mathcal{R}} B$. We want to show that (Result 2) holds when $|w| \vdash^{n+1}_{\mathcal{R}} B$. So let $n > 0$, and assume $\frac{1}{2}(B)$ at w and $|w| \vdash^{n+1}_{\mathcal{R}} B$. Then B has to be proved from certain assumptions $A_1, A_2, \ldots A_k$ such that each A_i is provable from $|w|$ in at most n steps and such that B is provable from $A_1, A_2, \ldots A_k$ in 1 step. There are three cases to consider:
>
>> **Case 1**: For some A_i, $0(A_i)$ at w. By the induction hypothesis, if $0(A_i)$ at w and $|w| \vdash^n_{\mathcal{R}} A_i$, then an i-contradiction ($\{1(A_i), 0(A_i)\}$) can be proved from $|w|$ in 1 step using $\mathcal{R}$.
>>
>> **Case 2**: For some A_i, $\frac{1}{2}(A_i)$ at w. By the induction hypothesis, if $\frac{1}{2}(A_i)$ at w and $|w| \vdash^n_{\mathcal{R}} A_i$, then an i-contradiction ($\{1(A_i), \frac{1}{2}(A_i)\}$) can be proved from $|w|$ in 1 step using $\mathcal{R}$.

[15] For the proof of (Result 2), I omit reference to the parts of Jago's formalism that concern the function f and the relation $[[\mathcal{R}]]$ between scenarios. Both simplifications are innocuous.

Case 3: For all A_i, $1(A_i)$ at w. We now repeat the argument from the base case. Since B can be derived from $A_1, A_2, \ldots, A_k$ in 1 step using $\mathcal{R}$, and since each $A_i \in |w|$, then $|w| \vdash^1_{\mathcal{R}} B$. Then B $1(B)$ is provable from $|w|$ in 1 step. Since neither $B \in |w|$ nor $\neg B \in |w|$, $\frac{1}{2}(B)$ at w. So an i-contradiction ($\{1(B), \frac{1}{2}(B)\}$) can be proved from $|w|$ in 1 step using $\mathcal{R}$.

Omitting the analogous reasoning for $0(B)$ at w, I conclude that (Result 2) holds for the inductive step and thus in general. ∎

(Result 2) states that an i-contradiction can be proved from any scenario that fails to be logically consistent in just 1 step of logical reasoning using $\mathcal{R}$. Effectively, (Result 2) shows that the attempt to construct a stratified epistemic space based on Jago's model will collapse when $\mathcal{J}$ satisfies (C_2). Since an i-contradiction can proved from any scenario that fails to be logically consistent in just 1 step, all scenarios that are not in W_0 will be logically consistent.

(Result 2) implies that we cannot use $\mathcal{J}$ to establish (Plentitude$_n$). For consider a sentence A that is disprovable in n, but not m steps of logical reasoning, for some m such that $n > m > 1$. Then we have $\Diamond_m A$. To establish (Plentitude$_n$), we need a scenario in W_m at which A is true. By (Result 2), however, there are no scenarios that fail to be logically consistent in any W_n, for $n \geq 1$. Since A can be true only at a logically inconsistent scenario, A hence cannot be true at any scenario in W_m. So $\Diamond_m A$ is true but there is no scenario in W_m at which A is true. Thus (Plentitude$_n$) fails.

Furthermore, (Result 2) implies (Omni$_2$):

(**Omni**$_2$) Any agent who can perform just 1 step of logical reasoning using the rules in $\mathcal{R}$ can a priori reject any scenario in $W^C \cup W^O$ that fails to be logically consistent.

Insofar as any minimally rational agent is capable of performing just 1 step of logical reasoning using the rules in $\mathcal{R}$, it follows from (Omni$_2$) that only logically consistent scenarios remain strictly epistemically possible for such agents. Derivatively, all logical truths are strictly epistemically necessary for all minimally rational agents, and, as such, they are wrongly characterized as logically omniscient.

This concludes my second argument. To avoid (Result 2), we cannot hold that some defective partial scenarios can be ruled out using rules in $\mathcal{R}$.

4.3 Argument against (C_3)

For the third argument, let us assume instead that $\mathcal{J}$ satisfies (C_3). Then scenarios in $\mathcal{J}$ can fail to be maximal but no defective partial scenarios can be ruled out using the rules in $\mathcal{R}$—although, of course, partial scenarios from which a contradiction of the form A and $\neg A$ is provable can still be ruled out. When $\mathcal{J}$ satisfies (C_3), epistemic space will contain an abundance of defective partial scenarios like $w^\star$ at which only "It rains and it is cold" is true but at which *all* other sentences are *indeterminate*.

Jago, I believe, does not assume that $\mathcal{J}$ satisfies (C_3). The main reason is that scenarios like $w^\star$ fail to capture "genuine epistemic possibilities for rational, yet non-ideal, agents." (Jago 2009, p. 340.) For, as motivated above, minimally rational agents always know that obvious logical falsities such as $(A \wedge \neg A)$ are false and that obvious logical truths such as $(A \rightarrow A)$ are true. To capture this central feature of non-trivial logical reasoning, we can hold that scenarios like $w^\star$—at which all obvious logical falsities and truths are indeterminate—should be rejected a priori by a minimally rational agent. If so, there is good reason to hold that $\mathcal{J}$ should not satisfy (C_3).

To illustrate this complaint in more detail, we can focus on deep$_n$ epistemic necessity—by doing so, we focus on the kinds of sentences that are provable within n steps of logical reasoning. Obvious logical truths can be proved within relatively few steps of logical reasoning, and by having an epistemic space that allows us to model deep$_n$ epistemic necessity, we can help ensure that all such truths are also strictly epistemically necessary for minimally rational agents. For deep$_n$ epistemic necessity, the task is to establish (N-Plentitude$_n$): $\Box_n A$ iff A is true at all scenarios w in W_n. If $\mathcal{J}$ satisfies (C_3), however, we cannot use $\mathcal{J}$ to ensure (N-Plentitude$_n$). For let A be an obvious logical truth that can be proved within, say, 4 steps of logical reasoning, and assume $\Box_4 A$. Then A should be true at all $w \in W_4$. But if very sparsely described scenarios like $w^\star$ are not eliminated from epistemic space, A may well fail to be true at a scenario in W_n, for *any* n. So $\Box_4 A$ is true but there are scenarios in W_4 at which A is *not* true. Thus (N-Plentitude$_n$) is false according to Jago's model. So if we want to capture (N-Plentitude$_n$), $\mathcal{J}$ cannot satisfy (C_3).

This concludes my third argument. If $\mathcal{J}$ satisfies (C_3), then $\mathcal{J}$ cannot serve as an epistemic space for minimally rational agents. To develop such a space, we need to eliminate scenarios that are *too* partially specified. Defective scenarios like $w^\star$ and countless others would fall into this category. The second argument, however, shows that attempts to eliminate some such partial scenarios fail because they end up eliminating all scenarios that are not logically consistent. So we have an argument showing that we cannot use $\mathcal{J}$ to ensure (Plentitude$_n$) and (N-Plentitude$_n$)—and hence an argument showing that we cannot use $\mathcal{J}$ to model strict epistemic possibility for minimally rational agents.

While the first argument shows that we must include partial scenarios in a suitable epistemic space for minimally rational agents, the third argument shows that we need to eliminate at least some (defective) partial scenarios from this space. In particular, we need to eliminate from epistemic space scenarios that fail to verify sentences that a minimally rational agent can easily establish logically (at a given time). Yet, if we specify a non-trivial method for eliminating such scenarios from epistemic space, the second argument shows that we thereby end up eliminating all scenarios that fail to be logically consistent. In other words, if there is a method for eliminating scenarios that fail to verify what follows by easy logical reasoning, this method will eliminate all scenarios whose truths are not closed under logical deduction simpliciter.

More generally, I believe, (Result 1) and (Result 2) show that (Information Principle) must go if we want to model strict epistemic possibility for minimally rational agents in a world-involving framework. For to model what remains strictly epistemically possible for such agents, we must eliminate from epistemic space scenarios that fail to contain information that can be inferred by easy logical reasoning. But if we begin to eliminate such scenarios from epistemic space, (Result 1) and (Result 2) show that we will end up eliminating all scenarios that fail to be closed under logical deduction simpliciter—and, in turn, that the resulting epistemic space will be useful only for modeling epistemic possibility for ideally rational agents. To avoid this result, we should not account for the gain in information that results from bounded, non-trivial logical reasoning in terms of *elimination* of scenarios. If so, (Result 1) and (Result 2) tell us that (Information Principle) must go in order to model strict epistemic possibility for minimally rational agents. And insofar as (Information Principle) is a constitutive principle of the possible worlds framework for knowledge and belief, we hence have reason to believe that we cannot use the framework to model strict epistemic possibility for minimally rational agents.[16]

[16]Notice that even if (Information Principle) is not strictly regarded as a constitutive principle of the possible worlds framework, it is certainly fair to say that it is among the core characteristic features of the framework. In that case, my results support the claim that we must give up a core feature of the possible worlds framework to model epistemic possibility for minimally rational agents.

5 Conclusion

I have argued that we cannot use Jago's model of epistemic space to model what is (strictly) epistemically possible for minimally rational agents. If all scenarios in $\mathcal{J}$ are maximal, it follows that minimally rational agents can rule out any logically inconsistent scenario in just 1 step of logical reasoning. If some scenarios in $\mathcal{J}$ are non-maximal, it follows either that minimally rational agents can rule out any scenario that fails to be logically consistent in just 1 step of logical reasoning, or that they can never rule out any defective partial scenarios. In the latter case, scenarios that do not describe a way the world might be, for all a minimally rational agent knows, must nonetheless survive in the corresponding epistemic space. These negative results support the claim that the prospects look dim for using the possible worlds framework to model what is (strictly) epistemically possible for minimally rational agents. In particular, as motivated, the results support the claim that we must give up (Information Principle) to develop an epistemic space that can model epistemic possibility for such agents.

While ordinary agents do indeed wonder about propositions that could not possibly obtain, they do not—at least not often under careful reflection—wonder about propositions that obviously could not obtain. For much like minimally rational agents, ordinary agents are able to tease out a few consequences of what they know, and to use their basic inferential capacities to establish that certain obvious truths such as "If it rains, then it rains" must hold true of our world. If the conclusions of this paper are correct, the prospects look dim for using the standard possible worlds framework to model what is epistemically possible for such ordinary agents.

References

Bjerring, J. C. (forthcoming). Impossible Worlds and Logical Omniscience: an Impossibility Result. Forthcoming in *Synthese*.

Buss, S. 1998. *Handbook of Proof Theory*. Vol. 137. Elsevier, Amsterdam.

Chalmers, D. J. (2011). The Nature of Epistemic Space. In A. Egan & B. Weatherson's (Eds.) *Epistemic Modality*. New York: Oxford University Press, pp. 60-107.

Cherniak, C. (1986). *Minimal Rationality*. Cambridge: MIT Press.

Drapkin, J. & Perlis, D. (1986). Step-logics: An Alternative Approach to Limited Reasoning. *Proceedings of the European Conference on Artificial Intelligence*. Brighton, England, pp. 160-163.

Field, H. (1977). Logic, Meaning and Conceptual Role. *Journal of Philosophy*, 69, pp. 379-408.

Field, H. (2001). *Truth and the Absence of Fact*. Oxford: Oxford University Press.

Harman, G. (1999). (Non-Solipsistic) Conceptual Role Semantics. In G. Harman's *Reasoning, Meaning, and Mind* . Oxford: Clarendon Press, pp. 206-231.

Hintikka, J. (1962). *Belief and Knowledge*. Ithaca: Cornell University Press.

Hintikka, J. (1975). Impossible Possible Worlds Vindicated. *Journal of Philosophical Logic*, 4, 475-484.

Hintikka, J. (2003). A Second Generation Epistemic Logic and Its General Significance. In V. Hendricks, K. Jörgensen, S. A. Pedersen (eds.), *Knowledge Contributors* (pp. 33-56), Synthese Library, vol. 322, Kluwer Academic Publishers, pp. 33-56.

Jago, M. (2009). Logical Information and Epistemic Space. *Synthese*, 167, pp. 327-341.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell Publishers.

Mendelson, E. (1997). *Introduction to Mathematical Logic*. 4th edition. London: Chapman & Hall.

Nolan, D. (1997). Impossible Worlds: A Modest Approach. *Notre Dame Journal of Formal Logic*, 38, pp. 535-572.

Priest, G. (2005). *Towards Non-Being: the Logic and Metaphysics of Intentionality.* Oxford: Oxford University Press.

Robbins, P. (2004). To Structure, or Not to Structure? *Synthese*, 139, pp. 55-80.

Stalnaker, R. (1984). *Inquiry.* Cambridge: MIT Press.

Wedgwood, R. (2007). *The Nature of Normativity.* Oxford: Clarendon Press.