Logical Information and Epistemic Space

Mark Jago (mark.jago@scmp.mq.edu.au)*
Department of Philosophy, Macquarie University, Sydney, Australia

Abstract. Gaining information can be modelled as a narrowing of *epistemic space*. Intuitively, becoming informed that such-and-such is the case rules out certain scenarios or would-be possibilities. Chalmers's account of epistemic space treats it as a space of a priori possibility and so has trouble in dealing with the information which we intuitively feel can be gained from logical inference. I propose a more inclusive notion of epistemic space, based on Priest's notion of *open worlds* yet which contains only those epistemic scenarios which are not obviously impossible. Whether something is obvious is not always a determinate matter and so the resulting picture is of an epistemic space with fuzzy boundaries.

Keywords: Logical information, epistemic space, epistemic scenario, epistemic possibility, impossible worlds, open worlds, rationality.

1. Introduction

Wittgenstein once remarked that 'there can *never* be surprises in logic' (1922, §6.1251), much to the surprise of anyone who's ever tried to master the subject. A cursory leaf through the technical literature shows that logic contains innumerable surprises. By way of example, students who happily accept the truth-table for the material conditional '$\rightarrow$' are surprised to learn that, for any propositions 'ϕ' and 'ψ' whatsoever, either $\phi \rightarrow \psi$ or $\psi \rightarrow \phi$. Even preeminent logicians are occasionally surprised by their results, as the reaction to early results in model theory shows. Löwenheim's theorem, stating that every satisfiable sentence (in a countable language) has a countable model, perplexed Skolem, who gave the first correct proof of it (Skolem, 1922). Skolem took the result to be paradoxical, for it entails that theories asserting the existence of an uncountable set, if consistent, have countable models. Dummett characterizes the situation well:

> When we contemplate the simplest basic forms of inference, the gap between recognizing the truth of the premises and that of the conclusion seems infinitesimal; but, when we contemplate the wealth and complexity of number-theoretic theorems, ... we are struck by the difficulty of establishing them and *the surprises they yield*. (Dummett, 1978, p. 297, my emphasis.)

* I would like to thank Greg Restall, Sebastian Sequoiah-Grayson, Timothy Williamson and the audience at Logica 2007 for useful suggestions on pervious versions of this paper.

A natural way to explain what makes a given truth surprising is to appeal to its informativeness. If logical results could not possibly be informative, then how can they be so surprising? One would most naturally characterize the import of Russell's famous letter to Frege of June 16, 1902 as *informing* Frege of the paradox of Basic Law V, via a logical deduction. And to take a practical example, we naturally say that a model checker will inform us whether a particular design meets a certain requirement.[1]

These brief remarks provide a prima facie case for the claim that at least some results in logic are informative. Dummett remarks that 'for [deduction] to be useful, a recognition of its truth need not actually have been accorded to the conclusion when it was accorded to the premises' (1978, p. 297). Deduction is useful because it allows us to recognize the (perhaps previously unrecognized) truth of the conclusion; but such recognition is itself a form of gaining information. Hintikka went so far as to call it a 'scandal' (1970, p. 289) that formal theories of information do not allow for deduction to be informative.[2] But the idea that deduction can be informative is met with resistance from a considerable number of logicians and philosophers. Logical reasoning is sometimes characterized as being 'trivial', not in the sense that it is always easy but in the sense that no new information results. Similarly, we often hear that deduction is the process of uncovering the information already contained in the premises. But there is no good reason for taking these characterizations literally.

Let us begin with the idea that deduction can only uncover information already contained within the premises. The thought might be that, if a valid deduction uncovers information not contained in the premises, then the conclusion drawn would tell us more about the world than the premises do. But, the worry continues, a valid conclusion cannot be a further truth about the world, one that is not guaranteed by the premises, for the premises of a valid argument must guarantee the truth of the conclusion. In holding that valid deductions can be informative, however, one is not claiming that such deductions reveal empirical facts which were not contained in or guaranteed by the premises. Information

[1] Model checking is a technique used extensively in industry as a way of verifying that certain properties, specified as logical formulae, hold of a system, specified as a theory. Model checking allows for flaws in a system to be discovered at the design stage with relative ease.

[2] Hintikka's own solution (1970; 1973), although ingenious, has the unwanted consequence that all inferences in the propositional and monadic first-order calculi are utterly uninformative. It is for this reason that Jago (2007) and Sequoiah-Grayson (2007) each reject Hintikka's proposal.

can be gained without a new fact being learnt (by presenting a fact in a new way, for example).[3]

In the *Tractatus* (1922), Wittgenstein held that tautologies lack any sense. They are *sinnlos*: they literally 'say nothing' (§6.11) and so 'theories which make a proposition of logic appear substantial are always false' (§6.111). His reason for thinking this is that meaningful propositions must form a 'logical picture' (§2.19) of the world, i.e., they must partition logical space into non-empty subspaces. Although this theory of meaning (and particularly the thesis that non-empirical propositions are meaningless) receives very little contemporary support, it does contain an idea which underlies a reason for holding that logical inference is uninformative. This is the idea of treating the concept in question—meaning, sense, or information—in terms of ways of partitioning logical space.

The idea is found in Carnap's (1947) notion of *intension* and a similar idea is found in Hintikka's (1962) groundbreaking work in modal epistemic logic. When applied to the concept of information, the key idea is that gaining information amounts to ruling out possibilities (of a particular kind).[4] But, as Wittgenstein saw, logical truths do not rule out any genuine possibility, so one might argue as follows. For a thought to be informative, it must be capable of narrowing down the set of admissible possibilities. Logical truths do not narrow down such possibilities and hence are uninformative (and similarly in the case of logical inferences). In what follows, I accept the first premise as a platitude but deny the second. Once the relevant notion of a possibility has been made precise, there is scope for logical truths to exclude at least some such possibilities, or so I shall argue.

The rest of the paper is organized as follows. In §2, I outline the concepts of an epistemic scenario and epistemic space and their relation to information. I discuss an alternative, non-idealized notion of epistemic space in §3, based on the notion of an *open world* and develop a modified version of this account, which incorporates our epistemic expectations of rational agents. I then develop formal models which capture this notion. Finally, in §4, I discuss how the resulting notion

[3] This, I take it, is the key insight in Frege's (1892) account of informative identity claims. (By 'fact', I mean an entity which does not contain Fregean senses.) If it were not the case that information can be gained by presenting a fact in a fresh way, for example, Jackson's knowledge argument against physicalism (1986) would be far more pressing. For then, when Mary first experiences redness, she must genuinely learn a new fact, which, given Jackson's premises, must be non-physical.

[4] Van Benthem (2003) gives a good overview of this approach to modelling information and its dynamics.

of an epistemic scenario admits borderline cases, before giving some concluding remarks in §5.

2. Ideal Epistemic Space

Let us call a way the world might be (irrespective of any a posteriori knowledge) an *epistemic scenario* (or, for brevity, just a *scenario*). Each scenario comes with a particular perspective, which can be treated as an individual and a time at the centre of that scenario. We can thus take scenarios to be (or to be modelled as) *centred worlds*: triples $\langle w, i, t \rangle$ containing a world plus an individual and a timepoint (Lewis, 1979, pp. 513–43).[5] Here, w is the world which underlies the scenario. Following Chalmers (2007), I will call the class of all epistemic scenarios *epistemic space*. The epistemic possibilities are then whatever is true according to some scenario in epistemic space (and dually, the epistemic necessities are whatever is true according to all such scenarios).

The information contained within a thought rules out certain scenarios as ways the world could be. So the informational content of a thought can be modelled as a function from scenarios to truth-values or, more simply, as a set of scenarios (assuming that *truth* and *falsity* behave classically).[6] In figurative terms, the informational content of a thought is a region of epistemic space. A thought is informative iff its informational content excludes at least some scenarios, for this gives rise to the possibility of some agent learning something new in grasping the thought. Similarly, a deductive process is informative iff engaging in it leads to at least some scenarios being ruled out as ways the world might be. Whether this picture of informational content can account for the apparent informativeness of some logical truths and some logical

[5] The individual and timepoint at the centre of a world allow for indexicals to be dealt with within each scenario. They also rule out scenarios in which no individual exists which, presumably, would not be an epistemic possibility for any agent. There are a number of issues surrounding the use of centred worlds. If one took an utterance of ‘I exist’ or ‘I am thinking’ to be a posteriori, for example, then there would have to be scenarios without an individual or without a thinker at its centre, respectively (see Chalmers, 2007, p. 11). For the purposes of the present paper, however, such issues can be placed to one side.

[6] Some, including Dretske (1981) and Floridi (2004), argue that only true thoughts can be informative, i.e., that information is factive. To ensure factivity, one need add only that, if ϕ is informative in world w, then the set of scenarios which models the informational content of ϕ must include w. Nothing I have to say here turns on whether information is factive and so I leave the issue open. (Strictly speaking, one should speak of a *doxastic* rather than an epistemic space when modelling a non-factive concept, but let’s not quibble.)

inferences depends on just what properties the worlds which underlie epistemic scenarios are assumed to have.[7]

Chalmers characterizes epistemic scenarios as 'maximally specific coherent ways the world might be' (2007, p. 1). 'Coherent' here means that no a priori falsehood is true according to any scenario. In particular, the truths according to such scenarios are closed under classical consequence so that, given their coherence, the set of truths according to any given scenario must have a classical model. The underlying world of each scenario, in other words, must be a logically possible world.[8] All a priori truths are true according to each epistemic scenario and no a priori falsehood is true according to any. As a consequence, the informational content of any a priori true thought is the entirety of epistemic space: in other words, a priori truths contain no information at all.[9] This view of epistemic space tells us that logical truths and valid inferences cannot possibly be informative.

This highlights a limitation inherent in the identification of epistemic scenarios with 'maximally specific coherent ways the world might be'. But this does not entail that one should reject epistemic space as a way of analysing information altogether. It is a platitude that the information contained within a thought allows an agent who accepts the thought as true to rule out certain scenarios. Because of this, the identification of that information with a region of epistemic space is highly intuitive: just as intuitive, in fact, as the thought that valid inferences can be informative. I take this to show that, rather than rejecting either intuition, one should first consider alternative notions of an epistemic scenario.

Before turning to discuss the alternatives, it should be noted that Chalmers is perfectly aware of these consequences for his notion of epistemic space. He remarks that 'the notion of epistemic possibility and necessity involves a rational idealization away from our contingent cognitive limitations' (2002, p. 612). As such, epistemic space is 'best suited for modeling the knowledge and belief of idealized reasoners'

[7] As there are epistemic possibilities which are not genuine metaphysical possibilities, for example, such worlds need not be metaphysically possible worlds (although see footnote 8).

[8] In fact, Chalmers argues that just the *metaphysically* possible worlds may underlie scenarios. This idea relies upon a distinction between a world *verifying* a sentence and that world *satisfying* it. According to Chalmers, some metaphysically possible worlds verify 'Hesperus is not Phosphorus', although none satisfy it (2007, p. 12). As this approach does not carry over into the construction of non-ideal epistemic space (§3 below), I will set it to one side.

[9] In this way, the view of what it is for a thought to be informative that we get from Chalmers's epistemic space comes close to the view of what it is for a proposition to be meaningful in the *Tractatus*.

(Chalmers, 2007, p. 8) and hence what is informative for an ideal agent, rather than what is informative for cognitively bounded agents such as you and I.[10] As ideal agents automatically know all of the valid sentences to be valid, there is no sense in which an ideal agent can discover anything new through valid reasoning.

Suppose we grant that this kind of idealization has a place.[11] One might then argue that, in contexts in which the idealization has been granted, it may be asserted that logical truths are uninformative. This would be a mistake, however. Whether a thought is informative or not depends on the senses which constitute it (and perhaps on whether it is true). For Frege, thoughts are ontologically independent of whether any agent grasps them and hence one cannot infer from the cognitive abilities of particular agents to the makeup of thoughts. One could account for the fact that ideal agents find logical truths uninformative by holding that they are unable to grasp any informative logically true thought.[12] This would in no way entail that informative, logically true thoughts do not exist. In short, one cannot justify the claim that logical truths are uninformative by appeal to the idealization of an agent's cognitive powers, because there is no route from idealization of such powers to idealization about the nature of thoughts themselves.

3. Non-Ideal Epistemic Space

A non-ideal epistemic space is one in which not all scenarios are maximally specific coherent ways the world might be. Some of these scenarios are not only metaphysically impossible but also impossible by the standards of classical logic: what is true according to such scenarios may be contradictory and need not be closed under classical consequence. Such scenarios are (or can be modelled using) centred *open worlds* (Priest, 2005). Open worlds are a species of impossible world, which

[10] Ideal agents are hypothetical agents from which all cognitive limitations (such as finite memory, time in which to reason and the like) have been abstracted.

[11] It is often made, harmlessly, when developing epistemic and doxastic logics, for there are many practical applications in which it does not matter whether agents are modelled as being logically omniscient. Perhaps it is this pragmatic kind of justification which Chalmers has in mind when he says that 'the idealized notion [of epistemic space] is the best-behaved and the easiest to work with' (2007, p. 8).

[12] Take a valid material equivalence, '$\phi \leftrightarrow \psi$', for example. For the corresponding thought to be informative, it must contain senses which present a common truth-value in distinct ways. For ideal agents, however, all logically equivalent ways of presenting a truth-value are uninformatively equivalent, so must constitute one and the same (ideal) sense. Hence thoughts containing distinct but logically equivalent senses of a truth-value cannot be grasped by an ideal agent.

allow each sentence to be true or false independently of the truth-value of any other sentence according to that world. Conjunctions and disjunctions may be true or false independently of the truth-value of their conjuncts and disjuncts, for example and '$\phi \wedge \psi$' may be true even if '$\psi \wedge \phi$' is false. So no inference rule (apart from the trivial inference from 'ϕ' to 'ϕ') is valid for open worlds in general.[13]

Treating epistemic scenarios as centred open worlds can, in principle, solve the problem for Chalmers discussed in the previous section. The worry is that the notion of epistemic space will become trivial if all centred open worlds are to count as scenarios. This is because there are open worlds according to which 0 is 1, squares are round, certain objects are both green all over and red all over and so on. If all such worlds are allowed to underlie epistemic scenarios, then *any* sentence whatsoever will describe an epistemic possibility. I shall argue, however, that these are not genuine epistemic possibilities for any agent.

Of course, it is not difficult to imagine an agent who professes to believe that there are (or might be) round squares or that 0 is (or might be) 1, or a cognitively impaired agent who cannot make the inference from '$a = b$' to '$b = a$'. One could then argue that round squares and the like are genuine epistemic possibilities for at least some agents. If there were such a notion of epistemic possibility, then all centred open worlds would count as epistemic scenarios and would be suitable for modelling the corresponding notion of information. The resulting picture would be of little philosophical or logical interest, for it takes informational content to be as highly structured as the syntax of the relevant language.[14] This notion bears little resemblance to our intuitive concept of information, to which we appeal when claiming that valid inferences can be informative.

One point which I take from Chalmers's discussion, therefore and which I take to be absolutely essential to any account of epistemic space, is that the concept of epistemic possibility should 'involve some imposition of a rational idealization' and so 'the corresponding notion of deep epistemic necessity should capture some sort of rational *must*' (Chalmers, 2007, p. 7). This point applies to epistemic concepts in general: we ascribe knowledge, belief, the ability to form judgements

[13] Formally, Priest treats each open sentence '$\phi(x_1, \ldots, x_n)$' at each open world on a par with a predicate, to which an extension (and, if one wants to allow for truth-value gluts and gaps, an anti-extension) is assigned. Then a closed sentence '$\phi(c_1, \ldots, c_n)$' is true according to an open world w iff the sequence of individuals denoted by 'c_1' through 'c_n' is a member of the extension assigned to '$\phi(x_1, \ldots, x_n)$' at w.

[14] Chalmers agrees, noting that the resulting notion of epistemic space 'will be of little use when it comes to analyzing meaning and content' (2007, p. 7).

and so on only to *agents*, i.e., systems which we take to be (or which can be interpreted as) rational to some degree. I do not, therefore, take the argument I gave at the end of the preceding section against Chalmers's notion of idealization to be an argument against any idealization being made when defining epistemic space. But I do depart company from Chalmers when he implicitly identifies 'rational idealization' with the end-point of ideal reasoning processes.

One can impose a criterion of rationality without thereby taking all agents to be ideal reasoners. A coherent, deductively closed set of beliefs is an ideal of rational enquiry, for example, yet an agent can be deemed rational if it has the ability to reason in accordance with certain logical rules and it deploys those abilities as well as the cognitive resources to hand allow. Failures of closure within an agent's belief set *may* be due to a failure of rationality but they may also be due to a lack of cognitive resources. One can, therefore, hold that the philosophically interesting notion of epistemic space is a rational space, incorporating the normative element of our epistemic concepts, without thereby holding it to be an ideal epistemic space in Chalmers's sense.

Our concept of epistemic possibility (and hence of an epistemic scenario) is a normative concept which, at the same time, should allow that valid inferences can be informative. In modelling this concept, the question that arises is, how should the notion of an epistemic scenario be constrained? Specifically, just which centred open worlds should count as epistemic scenarios? My suggestion is to focus on what is expected of a sincere, rational (although not ideal) agent. We expect such agents to recognize that there are no round squares, that 0 is 0 and not 1, that wholly green objects cannot be wholly red and so on. This notion of expectation plays an important role in our epistemic concepts; the difficulty is accounting for it in logical terms. For the remainder of this discussion, I will take it for granted that the normative component of our concept of epistemic possibility (and hence of an epistemic scenario) should be captured in terms of our expectations of rational agents (which, for brevity, I will refer to as our *epistemic expectations*). This notion can capture what Chalmers calls a 'rational *must*' (2007, p. 7), without leading to a level of idealization at which logical truths cannot be informative.

Chalmers does briefly consider non-ideal epistemic space and suggests that there may be a concept of epistemic possibility according to which ϕ is epistemically possible when it 'cannot be ruled out *through reasoning of a certain sort*' (2007, p. 34). In particular, he considers taking ϕ to be epistemically possible when it is not obvious a priori that $\neg\phi$. The approach I am proposing can be seen as fleshing out this proposal. In fact, I would suggest that our epistemic expectations are the more fundamental concept, for the reason why obvious a priori falsity is relevant to our concept of epistemic possibility is because we do not expect rational agents to entertain such beliefs.

Allowing every open world to underlie some epistemic scenario would clearly violate the 'rational must' requirement, for some open worlds give rise to truths which no agent can believe without falling below epistemic expectations. So I will take all of Chalmers's a priori coherent worlds, plus some but not all open worlds, to underlie the epistemic scenarios. More precisely, I will take a centred open world to be an epistemic scenario iff no agent would fall below our epistemic expectations by believing anything that is true according to that world.[15] As a consequence, some (but not all) pairs of jointly inconsistent sentences can be true according one and the same scenario and so truth according to such scenarios cannot be closed under classical consequence (for trivial worlds, according to which every sentence is true, will not underlie any scenario).[16]

To simplify matters, one may adopt the following scenario maximization principle. Let $|w|$ be the set of truths according to an open world w, i.e., w's *truth set.* The maximization principle requires that, if $|w| \subset |w'|$, then w is a scenario only if w' is not. This principle maximizes the truths that hold according to a consistent scenario. Accordingly, there will be just two kinds of epistemic scenario: those at which the truths are both deductively closed and consistent and those at which they are inconsistent (and hence not closed).

One way to capture formally (at least one aspect of) our epistemic expectations is to place a total order $\preceq$ on centred worlds, with $w \preceq w'$ when our expectations are such that, 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.[17] The maximal elements with respect to $\preceq$ are the worlds which are coherent (in Chalmers's sense):

[15] Our epistemic expectations can constrain the kinds of centres that scenarios may have. If such expectations require that no agent believes herself to be a teapot, then no scenario will have a teapot at its centre. Similarly, if we expect each agent to self-ascribe the ability to think, these expectations will force every scenario to have a thinker at its centre (see also footnote 5).

[16] Note that this also rules out inconsistent worlds at which truth is closed under paraconsistent consequence from underlying any scenario. Although not every sentence has to be true according to such worlds, some explicit contradictions will be true and holding an explicit contradiction to be true is one way of falling below epistemic expectations.

[17] In taking $\preceq$ to be a total order, I am assuming that any two scenarios may be compared with one another. If this turns out to be inappropriate, a partial order should be preferred; but I see no reason to think that any pair of scenarios are mutually incomparable.

for any such world w and all worlds w', $w' \preceq w$. The minimal elements with respect to $\preceq$ are those according to which some obvious a priori impossibility is true, where that impossibility is as basic as an a priori impossibility can be.

I take such basic a priori impossibilities to be those that any competent language user would recognize as false, non-inferentially, on a priori grounds. This includes identities of the form '$0 = 1$' and negated identities of the form '$a \neq a$', descriptions of round squares, objects that are both wholly green and wholly red and so on. The condition that their falsity be recognized non-inferentially excludes sentences such as 'if $1 = 1$, then a is both round and square'. As these basic impossibilities are captured in single sentences, I will allow that explicit contradictions (of the form '$\phi \wedge \neg\phi$') are recognized as falsehoods non-inferentially and so all such sentences are included as basic impossibilities.[18] Let x be a set of such basic a priori falsehoods. If anything at all is to be expected of a rational agent's beliefs, it is that no sentence in x is held as true. Given this notion, if $\mathsf{x} \cap |w| \neq \varnothing$ then $w \preceq w'$, for all worlds w'. Semantically, x can be treated as a set of centred open worlds $[\![\mathsf{x}]\!]$, such that $w \in [\![\mathsf{x}]\!]$ iff $|w| \cap \mathsf{x} \neq \varnothing$.

Next, we need some way to determine when $w \preceq w'$ holds for arbitrary centred worlds w and w'. The strategy I offer below is but one of perhaps many ways of doing this. The idea is as follows. In addition to expecting agents to treat each member of x as describing an epistemic impossibility, we also expect rational agents to perform basic inferences, in accordance with the meanings of 'and', 'or', 'for all' and so on. Note that it is one thing to have the ability to make such inferences but quite another to actually apply the corresponding inference rules to some fixpoint. So assuming that rational agents have a basic inferential ability is not to assume that their beliefs are closed under the corresponding rules. Let $\mathcal{R}$ be a set of basic inference rules, such that we can expect any rational agent to have the ability to apply any of those rules (if its cognitive resources allow). Semantically, $\mathcal{R}$ can be interpreted as a binary relation $[\![\mathcal{R}]\!]$ between worlds, where $(w, w') \in [\![\mathcal{R}]\!]$ iff there is a rule (instance):

$$\frac{\phi_1, \ldots, \phi_n}{\psi}$$

[18] It could be objected that one recognizes '$\phi \wedge \neg\phi$' to be false inferentially (using a one-step reductio). One could include pairs of sentences ('ϕ', '$\neg\phi$') in the set of basic impossibilities instead of such explicitly contradictory sentences. But as this would not affect the formalism greatly I will, merely for simplicity, stick with including all sentences of the form '$\phi \wedge \neg\phi$'.

in $\mathcal{R}$ such that $\{`\phi_1', \ldots, `\phi_n'\} \subseteq |w|$, $`\psi' \notin |w|$ and $|w'| = |w| \cup \{`\psi'\}$.

Let '$\Gamma \vdash_{\mathcal{R}} \phi$' abbreviate '$`\phi'$ is derivable from $\Gamma$ using just the rules in $\mathcal{R}$' and consider two sets of sentences, $\Gamma$ and $\Delta$. Suppose that $\Gamma \vdash_{\mathcal{R}} \phi$, $\Delta \vdash_{\mathcal{R}} \psi$, $\{`\phi', `\psi'\} \subseteq \mathsf{x}$ and that the minimum number of inference steps required to obtain 'ϕ' from Γ is less than the minimum number of steps required to obtain 'ψ' from Δ. Then there is an intuitive sense in which Γ is more obviously incoherent that Δ for, although both are incoherent, the incoherence of Δ is harder to spot (using $\mathcal{R}$) than the incoherence of Γ. If we expect rational agents to reject Δ as a description of an epistemic scenario, then our epistemic expectations require the agent to reject Γ as a description of a scenario too. Now, take two centred open worlds w and w' such that $|w| = \Gamma$ and $|w'| = \Delta$. It is in such cases that $w \prec w'$.

A model of epistemic space can then be defined with respect to a set of basic impossibilities x and a set of rules $\mathcal{R}$. The basic (propositional) model is a tuple:

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

where W^C and W^O are non-overlapping sets of centred closed and open worlds respectively such that, if $|w| \subset |w'|$ and $w \in W^O$, then $w' \notin W^C \cup W^O$.[19] $\preceq$ is a total order on $W^C \cup W^O$ and V is a propositional valuation function, assigning a truth-value to each sentence at each world.

To get the intended space, $\preceq$ must be constrained by x and $\mathcal{R}$. Define $[\![\mathsf{x}]\!] \subseteq W^C \cup W^O$ to be $\{w : |w| \cap \mathsf{x} \neq \varnothing\}$ and $[\![\mathcal{R}]\!]$ to be a binary relation on $W^C \cup W^O$, constrained by the rules in $\mathcal{R}$ as described above. Let f be a partial function from W^O to $\mathbb{N}$ such that $fw = n$ iff there is a sequence $w_0 w_1 \cdots w_n$ of centred worlds in $W^C \cup W^O$ such that $(w_i, w_{i+1}) \in [\![\mathcal{R}]\!]$ for each $i < n$, $w_0 = w$, $w_n \in [\![\mathsf{x}]\!]$ but no sequence $w_0 w_1 \cdots w_m$ with these properties is such that $m < n$. Finally, for any centred worlds w and w', set $w \preceq w'$ iff either $w' \in W^C - [\![\mathsf{x}]\!]$, $w \in [\![\mathsf{x}]\!]$, or $fw \leq fw'$.[20]

So defined, $\preceq$ captures at least part of the formal features of our epistemic expectations of rational agents. But such models do not yet say which centred worlds count as scenarios (some centred worlds in the model, including all those in $[\![\mathsf{x}]\!]$, will not count as scenarios). What can be gleaned from such models is that, if w is counted as a scenario and $w \preceq w'$, then w' too should be counted as a scenario. In the next section,

[19] By *closed world*, I mean a world w such that $|w|$ is closed with respect to classical logical consequence.

[20] I am assuming that, in the case of any closed world w, no rule in $\mathcal{R}$ allows a member of x which is not already a member of $|w|$ to be inferred from $|w|$. In other words, all worlds in $W^C - [\![\mathsf{x}]\!]$ are a priori coherent worlds and hence can be used to construct epistemic scenarios.

I will discuss how to obtain models which partition their domain of centred worlds into those that are scenarios and those that are not, that is, models which determine an epistemic space.

4. Borderline Epistemic Scenarios

One consequence of defining epistemic possibility in terms of our epistemic expectations is that it admits borderline cases. What can be expected of a rational agent's reasoning, just as what is obvious, is not always a determinate matter. There are clear epistemic impossibilities, including round squares, explicit contradictions and objects that are both completely green and completely red. There are also clear epistemic possibilities, including all a priori truths. But there are also borderline cases, including compound states of affairs that entail the existence of round squares but which do so in a way that is borderline obvious, so that whether our epistemic expectations require an agent to notice this incoherence is a vague matter. In this section, I discuss how this vagueness may arise and how it can be modelled.

Above, I followed Chalmers in taking an epistemic possibility to be whatever is true according to some epistemic scenario. This raises the following question: are there borderline epistemic possibilities in virtue of borderline truths at certain scenarios, or in virtue of what is true at some centred world which is a borderline scenario? First, note that a borderline epistemic possibility is not the same thing as an epistemically possible borderline case. As it is a posteriori whether Chalmers is hirsute or bald, it is *determinately* an epistemic possibility that he is borderline bald. Consequently, there exist epistemic scenarios in which Chalmers is borderline bald.

But now consider some borderline obvious impossibility ϕ, whose subject matter is nevertheless perfectly determinate, so that it is (determinately) obvious that 'ϕ' is either determinately true or determinately false. So there is no scenario in which 'ϕ' is a borderline case and yet whether ϕ is a genuine epistemic possibility is not a determinate matter. By definition, it is indeterminate whether an agent would fall beneath our epistemic expectations in believing that ϕ. Therefore, cases of borderline epistemic possibility must be dealt with by allowing for borderline epistemic scenarios, that is, for centred worlds which are neither determinately included in nor determinately excluded from epistemic space.

In principle, any of the standard approaches to vagueness could be employed to model epistemic space with borderline epistemic scenarios. Here, I briefly discuss how an epistemicist and a degrees-of-truth account might look. I will note a problem for each account but tentatively conclude in favour of a particular degrees-of-truth account which seems especially suited to modelling epistemic space with borderline scenarios. First note that, as $\preceq$ is a total ordering on worlds, it can be used to define a comparative notion of scenariohood, such that w can be said to be at least as scenario-like as w' when $w' \preceq w$. One can then assign a degree $\delta w \in [0, 1]$ to each centred world w, such that $\delta w = 1$ if $w \in W^C - [\![\mathsf{x}]\!]$, $\delta w = 0$ if $w \in [\![\mathsf{x}]\!]$ and $\delta w \leq \delta w'$ if $w \preceq w'$.[21] It is then natural to interpret δw as the degree of w's inclusion in epistemic space.

On the epistemicist view, however, δw will have a purely epistemic reading, in terms of the credence that a rational agent should assign to 'w is an epistemic scenario'. The epistemicist view is that there exists a sharp worldly boundary between scenarios and non-scenarios.[22] On this view, the vagueness inherent in scenariohood is the result of our ignorance of where the boundary falls and as such does not mandate a departure from classical semantics. For the epistemicist, a model of epistemic space (given an ordering $\preceq$ on centred worlds, as defined above) is simply a classical subset E of those centred worlds such that, if $w \in E$ and $w \preceq w'$, then $w' \in E$ as well. The informational content of a thought is then a sharp region of epistemic space. By including some but not all a priori incoherent centred worlds in epistemic space, such models allow some logical truths and some logical inferences to be informative.

Schiffer (2003) presents a problem for this view. Suppose that Jack is borderline wealthy, borderline old, borderline stylish, borderline bald and borderline clever. And suppose, in line with the epistemicist position, that we interpret our reluctance to assert whether Jack has or lacks these properties as a 0.5 credence in 'Jack is wealthy', 'Jack is old' and so on. Given elementary probability theory, our credence in 'Jack is wealthy, old, stylish, bald and clever' should be 0.03125; in other words, we should believe pretty strongly that Jack is not wealthy, old, stylish, bald and clever. But this seems wrong: I for one am not pretty certain that Jack is not wealthy, old, stylish, bald and clever (see also MacFarlane, 2008).

A natural alternative to the epistemicist view is to interpret δw as the truth-value of 'w is a scenario'. On this view, epistemic space has

[21] The restriction to $\mathcal{R}$ discussed in footnote 20 ensures that the first two conditions are mutually exclusive.

[22] This is not to deny that there are degrees of scenario-hood in a comparative sense (for even if 'tall' has a sharp extension, one tall person can be taller than another tall person).

fuzzy boundaries.[23] If epistemic space itself is fuzzy, then its subregions can be fuzzy (when they contain a centred world which is a borderline scenario). The informational content (qua region of epistemic space) of the thought that ϕ can be treated as a function $\mathcal{I}_\phi$ from centred worlds to $[0,1]$, defined as follows:[24]

$$\mathcal{I}_\phi w = \begin{cases} \delta w & \text{if `}\phi\text{'} \in |w| \\ 0 & \text{otherwise.} \end{cases}$$

Centred worlds according to which 'ϕ' is true are included in the informational content of the thought that ϕ to the degree that they are scenarios, whereas those according to which 'ϕ' is false are not included to any degree. On this account, as on the epistemicist interpretation, some logical truths and some valid inferences are informative, for coming to believe such logical truths and performing such inferences allows an agent to exclude some incoherent epistemic scenarios.

A prima facie reason for rejecting this degrees-of-truth account of epistemic space is the many problems associated with degrees-of-truth accounts of vagueness in general, including accounting for higher-order vagueness and giving a suitable semantics for the logical connectives.[25] A problem which appears particularly pressing for the degrees-of-truth account of epistemic space is that of comparisons between the informational contents of thoughts. Suppose (to take a simplistic example) that $\mathcal{I}_\phi$ and $\mathcal{I}_\psi$ agree on the value of all centred worlds except w and that $\mathcal{I}_\phi w = 0.01$ whereas $\mathcal{I}_\psi w = 0$.[26] ψ contains all the information that ϕ does and, in addition, excludes w as a possibility (which ϕ does not). But since w is just about a non-scenario, it is tempting to say that 'ϕ and ψ contain the same information' is much closer to truth than falsity. If the informational content of ϕ and ψ are given by $\mathcal{I}_\phi$ and $\mathcal{I}_\psi$ (which, after all, are perfectly classical functions), however, 'ϕ and ψ contain the same information' must be either determinately true or determinately false, for it is always a determinate matter whether $\mathcal{I}_\phi$ and $\mathcal{I}_\psi$ are the same function.

[23] It can be treated as a fuzzy set of centred worlds. In general, a fuzzy set X is a pair consisting of a classical set Y and a function from Y to $[0,1]$, known as the *membership function*, assigning to each member of Y a degree of membership in X. In the case of epistemic space, the underlying set is $W^C \cup W^O$ and the membership function is δ.

[24] The content can be seen as a fuzzy set of centred worlds, with $\mathcal{I}_\phi$ as its membership function and $W^C \cup W^O$ as its underlying set.

[25] It is often held, for example, that the Łukasiewicz semantics cannot be correct, as it allows tautologies to be less than fully true and contradictions to be less than fully false (see, e.g., Williamson, 1994).

[26] This could be the case if $\delta w = 0.01$, 'ϕ' $\in |w|$ but 'ψ' $\notin |w|$.

Two responses are available to this worry. The first is to concede the objection but hold that fuzzy sets of centred worlds are merely models of information and, as such, do not allow one to infer that ϕ and ψ differ in their informational content from the fact that $\mathcal{I}_\phi$ and $\mathcal{I}_\psi$ are distinct. One who takes this line can point out that, quite generally, it is not the case that all properties of a model are properties of what is modelled. Motivated by the idea which drives the objection, one could then define the degree of truth of 'ϕ and ψ contain the same information' to be the maximal degree of any centred world included in the content of one of the thoughts (to some extent) but not the other (to any extent), i.e.:

$$\max\{\max\{\mathcal{I}_\phi w, \mathcal{I}_\psi w\} - \min\{\mathcal{I}_\phi w, \mathcal{I}_\psi w\} \mid w \in W^C \cup W^O\}.$$ [27]

Following this idea, one could also take the degree of truth of 'ϕ is informative' to be given by substituting δ for $\mathcal{I}_\psi$ in the above calculation.[28] This response to the objection is unattractive, however, for it concedes that such models of information are at best rough approximations to the concept. One would expect that, given a *good* model of the informational content of ϕ and ψ, the modelled content of each will coincide with that of the other just in case ϕ and ψ in fact contain the same information.

The second response available to a defender of the degrees-of-truth account of epistemic space is to reject the intuition behind the original objection and to maintain that statements such as 'ϕ and ψ contain the same information' are always determinately true or determinately false. This is the position I prefer. On the face of it, however, this would appear to concede the epistemicist's point, as there would then exist perfectly sharp boundaries, e.g., between centred worlds which are epistemic scenarios to degree n and all other worlds. There is, however, a variant of the degrees-of-truth account which is perfectly consistent with the existence of sharp, unknowable boundaries, to which I now turn briefly.

Consider again the problem which Schiffer raised for epistemicist theories, which suggests that the partial beliefs associated with borderline cases are not the results of uncertainty. With an uncertainty-related partial belief, one is generally not in the best epistemic position to make a judgement, yet one may hold a vagueness-related partial belief, e.g. that Jack is rich, and yet be in the best possible epistemic situation. MacFarlane (2008) proposes that, in the way in which belief can be thought of as taking-to-be-true, vagueness-related partial

[27] Here, $\max\{\mathcal{I}_\phi w, \mathcal{I}_\psi w\}$ corresponds to the fuzzy union of $\mathcal{I}_\phi w$ and $\mathcal{I}_\psi w$, whereas $\min\{\mathcal{I}_\phi w, \mathcal{I}_\psi w\}$ corresponds to their intersection.

[28] This simplifies to $\max\{\delta w - \mathcal{I}_\phi w \mid w \in W^C \cup W^O\}$, for $\max\{\mathcal{I}_\phi w, \delta w\} = \delta w$ and $\min\{\mathcal{I}_\phi w, \delta w\} = \mathcal{I}_\phi w$ for every world w.

beliefs should be thought of as taking-to-be-partially-true. If truth is the aim of belief, then partial truth (e.g., truth to degree 0.5) is the aim of vagueness-related partial belief (2008, pp. 13–14). This move, according to MacFarlane, can explain attitudes of partial endorsement without succumbing to Schiffer's objection. MacFarlane takes this to be a motivation for degree theory which is independent of the usual arguments from the sorites and from comparatives; hence, he says, a degree theory can be accepted whilst also accepting hidden semantic boundaries (2008, p. 24).

MacFarlane calls the position which does just that 'fuzzy epistemicism' and argues that it defuses the traditional problems for degree theories. Higher-order vagueness, for example, is handled by fuzzy epistemicism in much the same way it is handled by epistemic theories, by appeal to sharp, unknowable boundaries (although these are boundaries between degrees of truth, not between truth and falsity). Fuzzy epistemicism provides a neat solution to the problem posed for the degrees account of epistemic space discussed above. If $\mathcal{I}_\phi$ and $\mathcal{I}_\psi$ differ on the value of any world, then 'ϕ and ψ contain the same information' is determinately false, as the account above implies. But this fact may be beyond anyone's ken. Accepting both degrees of truth and hidden semantic boundaries is just what is required to motivate the degrees account of epistemic space and reject the intuition that comparisons between the informational contents of thoughts cannot always be a determinate matter.

5. Conclusion

This paper was motivated by two intuitions: firstly, that logical truths and logical inferences can be informative and, secondly, that the informational content of a thought should be captured in terms of the epistemic scenarios it rules out. Together, these require a non-ideal notion of an epistemic scenario. Yet, since the notion of a scenario is related to that of epistemic possibility, it is in part a normative notion. This is why I rejected the simplistic amendment to Chalmers's notion of epistemic space, discussed at the beginning of section 3, which allows any centred open world to count as an epistemic scenario.

I argued that an agent can be rational even if its beliefs fall below the stringent standards of an idealized reasoner. Rather, we judge an agent to be rational against the background of our epistemic expectations. This notion, it seems to me, introduces an appropriate level of normativity into the picture. I then suggested a model which interprets our epistemic expectations as an ordering on centred worlds. The resulting notion of an epistemic scenario is more inclusive than Chalmers's idealized notion yet more selective than the near-trivial notion according to which (just about) any set of sentences is the truth set of some scenario.

A consequence is that the notion of an epistemic scenario admits of borderline cases. Just how one fixes the extension of 'is a scenario' in a model will depend on one's view of vagueness in general, although I suggested some reasons for preferring a degrees-of-truth account, based on arguments from Schiffer and MacFarlane. On this view, the informational content of a thought is a (possibly fuzzy) region of epistemic space. Logical truths can be informative, for the contents of some logical truths exclude some epistemic scenarios. Although such excluded scenarios are incoherent, they are subtly incoherent and so are not trivial or obvious impossibilities. This is why they correspond to genuine epistemic possibilities for rational, yet non-ideal, agents.

There is much more to be said about how our epistemic expectations determine normative, rational but non-ideal notions of epistemic space and hence of informational content. One aspect not discussed here is how such expectations affect the way in which scenarios must be centred (see footnotes 5 and 15). It would also be interesting to investigate alternative ways of capturing formally our epistemic expectations and to compare the resulting models of epistemic space to the ones presented here. Perhaps the most interesting aspect of epistemic space discussed here is the notion of borderline epistemic scenarios. It would be interesting to investigate the effect of using a fuzzy, non-ideal epistemic space to model other epistemic concepts, such as knowledge and belief, which are usually dealt with in terms of ideal epistemic space. And, as Chalmers notes, there might be interesting relationships between rational yet non-ideal epistemic space and Fregean senses. But I leave these topics for future work.

References

Carnap, R.: 1947, *Meaning and Necessity.* University of Chicago Press.

Chalmers, D.: 2002, 'The Components of Content'. In: D. Chalmers (ed.): *Philosophy of Mind: Classical and Contemporary Readings.* Oxford University Press, pp. 608–33.

Chalmers, D.: 2007, 'The Nature of Epistemic Space'. `http://consc.net/papers/espace.pdf`, accessed 27/10/07.

Dretske, F.: 1981, *Knowledge and the Flow of Information.* Oxford: Blackwell.

Dummett, M.: 1978, 'The Justification of Deduction'. In *Truth and Other Enigmas.* Cambridge, MA: Harvard University Press, pp. 166–85.

Floridi, L.: 2004, 'Outline of a Theory of Strongly Semantic Information'. *Minds and Machines* **14**(2), 197–222.

Frege, G.: 1892, 'Über Sinn und Bedeutung'. *Zeitschrift für Philosophie und Philosophische Kritik* **100**, 25–50.

Hintikka, J.: 1962, *Knowledge and Belief.* Ithaca, N.Y.: Cornell University Press.

Hintikka, J.: 1970, 'Surface Information and Depth Information'. In J. Hintikka and P. Suppes (eds.): *Information and Inference.* Dordrecht: D. Reidel.

Hintikka, J.: 1973, *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic.* Oxford: Clarendon Press.

Jackson, F.: 1986, 'What Mary Didn't Know'. *The Journal of Philosophy* **83**(5), 291–95.

Jago, M.: 2007, 'Hintikka and Cresswell on Logical Omniscience'. *Logic and Logical Philosophy* **15**(4), 325–54.

Lewis, D.: 1979, 'Attitudes *De Dicto* and *De Se*'. *The Philosophical Review* **88**(4), 513–43.

MacFarlane, J.: 2008, 'Fuzzy Epistemicism'. http://johnmacfarlane.net/fuzzy-epistemicism.pdf (draft version), accessed 28/01/08.

Priest, G.: 2005, *Towards Non-Being.* Oxford: Clarendon Press.

Schiffer, S.: 2003, *The Things We Mean.* Oxford: Oxford University Press.

Sequoiah-Grayson, S.: 2007, 'The Scandal of Deduction'. *Journal of Philosophical Logic*, `www.springerlink.com`, DOI: 10.1007/s10992-007-9060-4.

Skolem, T.: 1922, 'Einige Bemerkungen zur Axiomatischen Begründung der Mengenlehre'. In: *Selected Works in Logic.* Oslo: Universitetsforlaget, pp. 137–52.

Van Benthem, J.: 2003, 'Logic and the Dynamics of Information'. *Minds and Machines* **13**(4), 503–19.

Williamson, T.: 1994, *Vagueness.* London: Routledge.

Wittgenstein, L.: 1922, *Tractatus Logico-Philosophicus.* London: Routledge & Kegan Paul.