The Content of Deduction

MARK JAGO Forthcoming in *Journal of Philosophical Logic*. Draft of November 2011.

Abstract: For deductive reasoning to be justified, it must be guaranteed to preserve truth from premises to conclusion; and for it to be useful to us, it must be capable of informing us of something. How can we capture this notion of information content, whilst respecting the fact that the content of the premises, if true, already secures the truth of the conclusion? This is the problem I address here. I begin by considering and rejecting several accounts of informational content. I then develop an account on which informational contents are indeterminate in their membership. This allows there to be cases in which it is indeterminate whether a given deduction is informative. Nevertheless, on the picture I present, there are determinate cases of informative (and determinate cases of uninformative) inferences. I argue that the model I offer is the best way for an account of content to respect the meaning of the logical constants and the inference rules associated with them without collapsing into a classical picture of content, unable to account for informative deductive inferences.

Keywords: Content, information, deduction, inference, epistemic scenarios.

1 Introduction

DEDUCTIVE REASONING is essential to philosophy, mathematics and logic. What is not so clear is how deductive reasoning conveys information to us. In 'The Justification of Deduction', Dummett (1978, 297) asks how deduction can be both justified and useful. If it is justified, it must be guaranteed to preserve truth from premises to conclusion. To be useful, it must inform us of something.[1] But how, wonders Dummett, can the move from premises to conclusion be informative, if the former already guarantee the latter?[2] The task is to capture this notion of information content whilst respecting the fact that the content of the premises, if true, already secures the truth of the conclusion. This is the aim of this paper.

According to a popular analysis, for a proposition to be informative is for it to rule out certain scenarios, or would-be possibilities.[3] The proposition *that it rarely rains in Cambridge* is informative because it excludes possible scenarios in which it rains regularly in Cambridge. Before coming to believe that proposition, it was possible as far as I was concerned that it rains regularly in Cambridge. In coming to believe that proposition, I cease to treat such scenarios as ways the world might

1. 'For [deduction] to be legitimate, the process of recognising the premisses as true must already have accomplished whatever is needed for the recognition of the truth of the conclusion; for it to be useful, a recognition of its truth need not actually have been accorded to the conclusion when it was accorded to the premisses' (Dummett 1978, 297).

2. As Dummett notes, 'no definite contradiction stands in the way of satisfying these two requirements'; yet 'it is a delicate matter so to describe the connection between premisses and conclusion as to display clearly the way in which both requirements are fulfilled' (Dummett 1978, 297).

3. Hintikka (1962) gives the classic presentation of the view; Lewis (1975; 1986), Stalnaker (1976; 1984) and Chalmers (2002; 2010) put it to work in various ways. For recent overviews of the literature, see van Benthem and Martinez 2008 and van Benthem 2011.

be, for all I know. We can think of all scenarios according to which it rarely rains in Cambridge as constituting a notion of content for 'it rarely rains in Cambridge' which is suitable for various epistemic purposes.[4] To believe that proposition is to treat no other scenario as being a doxastic possibility; to know that proposition is to treat no other scenario as being an epistemic possibility. That content is informative for an agent iff coming to believe (or know) that proposition narrows down her doxastically (or epistemically) accessible scenarios. To be informative at all, therefore, a statement must have a non-empty content.

How should we think of the content of a valid deduction $\Gamma \vdash A$, from premises Γ to conclusion A, within this framework?[5] We can think in terms of the differences in an agent's belief state in the move from believing each premise in Γ but not believing the conclusion A to believing A as well as all premises in Γ. Alternatively, we can think in terms of an agent discovering the incompatibility of the premises Γ with the *negated* conclusion, $\neg A$. I'll write '$|A|$' to denote the set of all scenarios which represent that A is true, and '$|\Gamma|$' for the set $\bigcap_{A \in \Gamma} |A|$, i.e. the set of all scenarios which represent each member of Γ as being true.[6] On the first view, the content of a deduction $\Gamma \vdash A$ comes out as $|\Gamma| - |A|$ (or equivalently, $|\Gamma| \cap (|A|)^c$): the set of all scenarios which verify each member of Γ, but do not verify A.[7] Call this notion *content*$_1$. On the second view, the content of $\Gamma \vdash A$ comes out as $|\Gamma| \cap |\neg A|$: the set of all scenarios which verify each member of Γ and also the negation of A. Call this notion *content*$_2$. When negation behaves classically at all scenarios (i.e. when each s verifies $\neg A$ just in case it does not verify A) these two notions of content are equivalent: the set of *logically possible* scenarios which verify Γ but not A is precisely the set of those which verify both Γ and $\neg A$.

Our problem is then that, on either notion of content, there are no such (classical) logically possible scenarios whenever A can be derived logically from Γ. Hence each notion of content is empty, and the valid deduction $\Gamma \vdash A$ is treated as utterly uninformative. To resolve the puzzle, we require, at a minimum, a notion of a scenario which allows scenarios which verify Γ and either verify $\neg A$ or fail to verify A, even when Γ entails A. I'll assume throughout that we are interested in a classical notion of entailment (so that 'valid deduction' will mean a *classically* valid one). If we want to give a non-empty content to such deductions, therefore, we will need to appeal to non-classical scenarios.[8] This task is relatively

4. There are many other notions of content available. Throughout, I will take 'content' to mean the specific epistemic notion I've just described.

5. I won't make anything of the distinction between a specific *procedure* of deducing A from Γ and the statement that A can be deduced from Γ, '$\Gamma \vdash A$'. I'll talk throughout in terms of the content of a deduction $\Gamma \vdash A$. In saying that $\Gamma \vdash A$ is informative, I mean that performing some derivation of A from Γ has the potential to be informative to some agent (or: it would be informative to an agent who begins with no declarative information at all). Thus, $\Gamma \vdash A$ is informative only if its content is non-empty.

6. Throughout, A, B, C etc. are names for sentences. So I talk about a scenario representing that A is true, rather than representing that A. I'll also say that a scenario *s verifies A*, merely to abbreviate 's represents that A is true'. I'll use the convention that all logical formulae come with implicit Quine-quotes. Thus $\neg A$ abbreviates $\ulcorner \neg A \urcorner$.

7. Here, $(|A|)^c$ is the set-theoretic compliment of $|A|$ on the set of all scenarios.

8. Once we admit non-classical scenarios, content$_1$ and content$_2$ may come apart, for there may be scenarios which verify neither A nor $\neg A$ or both A and $\neg A$ (although see §2 for an argument that

easy, for many such structures are well-understood in contemporary model theory. The harder task is to explain why such structures should count as (or are good representations of) epistemically or doxastically possible scenarios. To provide a suitably epistemic notion of content, in terms of which we can say what it is for a deduction to be informative, the scenarios we use must all represent epistemic (or doxastic) possibilities: they must represent what *seem* to be genuine possibilities.[9]

The remainder of the paper proceeds as follows. In §§2–3 I present and discuss two accounts of non-classical scenarios which, it has been suggested, are suitable for our purposes. I argue that neither notion is suitable for our purposes. In §4, I discuss an analogy between our problem and the sorites paradox, concluding that the notion of content we are seeking is inherently vague (and as such, we should reject attempts to model contents by imposing artificial precision). Then, in §5, I present a formal model of content which treats some (but not all) valid deductions as being informative. I evaluate the model in §6.

2 The Relational Models Approach

To resolve our problem, we require scenarios which may verify both Γ and $\neg A$, or alternatively verify Γ but not verify A, even when $\Gamma \vdash A$. A popular place to look for such a notion is the model theory of *paraconsistent* and *paracomplete* logics, in which negation behaves in non-classical ways.[10] In a model of a paraconsistent logic, a sentence may be both true and false; in a model of a paracomplete logic, a sentence may be neither true nor false. In this section, I'll evaluate the suggestion that such models provide a way to model the content of deduction.

I'll call such models *relational models*, because we obtain them by replacing the classical valuation function with a *relation* V between models, sentences and the classical truth-values, so that a model may relate a sentence to either, both or neither of the values 1 and 0 (standing for *truth* and *falsity*).[11] Truth-at-a-model ($\vDash$) is then defined as follows:

---

there are no genuinely epistemic scenarios of the latter kind).

9. Note that allowing logically impossible scenarios to play an *epistemic* role does not entail that one can know impossibilities (e.g., true contradictions). What one knows is a matter of what is the case in all accessible epistemic scenarios. Standardly in epistemic logic, we count the actual world as an epistemic scenario for all agents. Then, if there are no actual true contradictions, no one will be modelled as knowing any true contradiction.

10. Priest (1987; 2008) gives the background to these logics. Models of these logics have a history in the epistemic logic literature, particularly in connection with worries about logical omniscience (Levesque 1984; Lakemeyer 1986; 1987; 1990).

11. There are alternative model theories for paraconsistent and paracomplete logics, which I won't go into here.

$s \vDash p$	iff	$Vsp1$
$s \vDash \neg p$	iff	$Vsp0$
$s \vDash \neg\neg A$	iff	$s \vDash A$
$s \vDash A \wedge B$	iff	$s \vDash A$ & $s \vDash B$
$s \vDash \neg(A \wedge B)$	iff	$s \vDash \neg A$ or $s \vDash \neg B$
$s \vDash A \vee B$	iff	$s \vDash A$ or $s \vDash B$
$s \vDash \neg(A \vee B)$	iff	$s \vDash \neg A$ & $s \vDash \neg B$
$s \vDash A \to B$	iff	$s \vDash \neg A$ or $s \vDash B$
$s \vDash \neg(A \to B)$	iff	$s \vDash A$ & $s \vDash \neg B$

We use relational models to give an account of the content of a classically valid deduction by allowing relational models to count as scenarios. To see how this helps, let's focus on a classical deduction involving some *modus ponens* steps, from $A \to B$ and A to B. We can find relational models where both premises are true but B is not. Suppose A is both true and false, and B is (just) false, at s. Then we have $s \vDash A$ and $s \vDash \neg A$ and hence $s \vDash A \to B$, but not $s \vDash B$. Thus the content$_1$ of the inference from $A \to B$ and A to B, defined as $(|A \to B| \cap |A|) - |B|$, is non-empty. We also have $s \vDash \neg B$, and so the content$_2$ of the inference, defined as $|A \to B| \cap |A| \cap |\neg B|$, is also non-empty. In this way, by allowing relational models to count as scenarios, we can provide a non-empty content$_1$ and content$_2$ for some classically valid deductions.

On this picture, not all valid deductions come out as being contentful$_1$. The deduction $A, B \vdash A \wedge B$ remains contentless$_1$, since any scenario verifying both A and B also verifies $A \wedge B$ (and so $|\{A, B\}| - |A \wedge B|$ is empty). Indeed, any classically valid inference not involving '$\neg$' or '$\to$' will be deemed contentless$_1$, on this view. This is a bad consequence for an account of content. *Modus ponens* and conjunction elimination, for example, do not seem to be wholly different kinds of inference rule. If one is deemed trivial or obvious, then the other should be too. One kind of inference should not be deemed to have content if the other does not.

Does our notion of content$_2$, defined as $|\Gamma| \cap |\neg A|$ for premises Γ and conclusion A, fare any better? It is easy to see that, when the set of scenarios includes relational models, *every* valid inference is deemed contentful$_2$. Take our conjunction introduction case from above (which was not deemed contentful$_1$). If scenario s verifies both A and B and in addition verifies either $\neg A$ or $\neg B$, then it will also verify both $A \wedge B$ and $\neg(A \wedge B)$. Such scenarios comprise the set $|A| \cap |B| \cap |\neg(A \wedge B)|$, and hence constitute the content$_2$ of $A, B \vdash A \wedge B$. Even the most trivial inference of all, from A to A, is deemed contentful$_2$: $|A| \cap |\neg A|$ is the set of all 'glutty' scenarios which verify both A and $\neg A$. Since any sentence A whatsoever can be represented as being both true and false by a scenario, $|\Gamma| \cap |\neg A|$ is guaranteed to be non-empty, regardless of any deductive relationships between Γ and A. This too is a bad consequence for a notion of content. I do not know of any reason for taking $A \vdash A$ (for example) to be informative. So, on either notion of content, the relational models approach does not provide a good account of the content of a valid deduction.

There is a deeper problem with the relational models approach: it fails to explain why the scenarios it provides are suitable tools for analysing epistemic notions of content and information. It is a consequence of the account that both the content_1 and content_2 of a valid deduction $\Gamma \vdash A$ can contain only glutty models, which assign both 0 and 1 to some sentence.[12] Yet it is a constraint on epistemic scenario-hood that what a scenario represents as being the case must at least seem possible. We want to model an agent's learning that A in terms of her ruling out some scenarios. If what those scenarios represent as being the case is obviously impossible to any agent who meets minimal standards of rationality, then there's no sense in which ruling out those scenarios corresponds to gaining new information.

From a classical point-of-view, it is obviously, trivially impossible for both A and $\neg A$ to be true. Dialethists will disagree, of course; and such debates can't be settled here. We can agree to fix our attention on agents who take it to be obviously, trivially impossible for both A and $\neg A$ to be true.[13] It might even be partially constitutive of their notion of being a rational agent that one takes it to be obviously, trivially impossible for both A and $\neg A$ to be true. We can build this principle into our notion of informativeness *for those agents*, so that any representation of both A and $\neg A$ is ruled out from the start.[14] Consequentially, in modelling what's informative for those agents, we shouldn't treat such representations as scenarios at all.[15] Once we slim down our class of (classical plus relational) scenarios in this way, we throw out all the relational models and are left only with classical models. As we have already seen, none of these scenarios allows us to provide a non-empty content_1 or content_2 for any classically valid inference. Expanding our stock of scenarios with relational models has not helped us to provide non-empty contents for classically valid inferences.

This is the very feature which makes our problem difficult. If we are to model the content of a valid deduction as a set of scenarios, then we have to admit 'impossible' (non-classical) scenarios. But trivially impossible models of explicit contradictions cannot feature in any account of rational (but non-ideal) attitudes. And on the relational models account of deductive content, the trivially impossible

12. To see why, assume that $\Gamma \vdash A$. Then on any consistent truth-assignment to the non-logical vocabulary of $\Gamma \cup \{A\}$ on which all members of Γ are assigned 1, A will be assigned 1 as well. (This holds regardless of whether the truth-assignment as a whole is classical or paracomplete, i.e., allowing that some sentences are assigned neither 0 nor 1.) No scenario in $|\Gamma| - |A|$ or $|\Gamma| \cap |\neg A|$ corresponds to such a truth-assignment. Hence all scenarios in $|\Gamma| - |A|$ and $|\Gamma| \cap |\neg A|$ (content_1 and content_2, respectively) assign both 0 and 1 to some sentence in the non-logical vocabulary of $\Gamma \cup \{A\}$.

13. Indeed, we could even work with agents who allow that it's *possible* for both A and $\neg A$ to be true, but for whom it is obviously, trivially not in fact the case that both A and $\neg A$ are true.

14. Of course, *obviousness* is both highly subjective and a matter of degree. This does not affect the argument, since an explicit contradiction $A, \neg A$ represents a determinate case of an obvious impossibility for all for the agents under consideration.

15. We must be precise when stating this objection. It is *not* that the relational models approach is incompatible with a rational agent's belief in each instance of $\neg(A \wedge \neg A)$. For this is paraconsistently equivalent to $A \vee \neg A$, which is verified by any scenario which assigns some value to A. The objection is not that the approach conflicts with what rational agents do in fact believe. Rather, the objection is that, given the uses to which scenarios will be put, the very notion of a scenario excludes such obviously glutty models.

models are *all* we're left with. In short, our question is difficult because it requires us to find scenarios which are impossible, but not trivially so.

In this section, I've argued that, when we interpret scenarios as relational models, neither content_1 nor content_2 provides a satisfying account of how a valid deduction can be informative. I then argued that (at least with respect to a certain background notion of rationality) we shouldn't count any model of an explicit contradiction as an epistemic scenario. If so, then both the content_1 and content_2 of any valid deduction comes out empty, on the relational models approach. In the next section, I consider an attempt to make sense of models which are impossible, but subtly so.

3 The Urn Model Approach

In this section, I discuss an attempt by Hintikka (1975) and Rantala (1975) to provide formal models which 'look possible but which contain hidden contradictions' (Hintikka 1975, 476). They aim to characterise models 'so subtly inconsistent that the inconsistency could not be known (perceived) by an everyday logician, however competent' (Hintikka 1975, 478). As we saw in the previous section, relational models are not appropriate in an account of epistemic content, precisely because whenever they represent the impossible, they represent the trivially impossible. If Hintikka and Rantala are successful in their aim, then we have good candidates to play the role of scenarios in our account of content.

The models Hintikka and Rantala describe are based on Hintikka's (1973a; 1973b) game-theoretic models of classical first-order logic.[16] The world is viewed as an urn from which individuals are drawn by two players, called '$\forall$' and '$\exists$'. In a game $G(\forall x A)$, player $\forall$ must pick an individual from the urn satisfying A; if she picks individual a, the game continues as $G(A[a/x])$. Similarly, the game $G(\exists x A)$ requires $\exists$ to pick an individual a satisfying A and continues as $G(A[a/x])$.[17] In this way, nested quantifiers represent constraints on sequences of draws from the urn. Just as in probability theory, individuals can but need not be replaced after being drawn from the urn. Models in which individuals are always replaced immediately after being drawn are the *invariant* models; all others are *changing* models. Invariant models provide a classical first-order semantics, whereas changing models are non-classical (and correspond to what Hintikka (1975) terms the *exclusive* interpretation of the quantifiers).

Hintikka's idea is that, given a sentence of certain game-theoretic complexity, there is a set of changing models 'which vary so subtly as to be indistinguishable from invariant [i.e., classical] ones at a certain level of logical analysis' (Hintikka 1975, 483). Hintikka (somewhat unfortunately) calls such models 'impossible possible worlds'. For the formal details, Hintikka draws on Rantala's (1975) notion of an *urn model*:

16. Hintikka's models develop the ideas of Henkin (1961) and Peirce (1992).

17. Analogously, $\forall$ decides whether $G(A \wedge B)$ proceeds as $G(A)$ or as $G(B)$ whereas $\exists$ decides how $G(A \vee B)$ should proceed. The game $G(\neg A)$ proceeds as the inverse game $\overline{G(A)}$, in which the players swap roles.

Definition 1 (Urn sequence) *An urn sequence* Δ *over a domain* $\mathcal{D}$ *is a countable sequence* $\langle D_i \mid i \in \mathbb{N}\rangle$ *where* $D_1 = \mathcal{D}$ *and, for* $i \geq 1$*,* $D_i \subseteq \mathcal{D}^i$ *(the ith Cartesian power of the domain* $\mathcal{D}$*) such that, for some* $a' \in \mathcal{D}$*:*

$$\langle a_1 \cdots a_i \rangle \in D_i \ \text{ only if } \ \langle a_1 \cdots a_i a' \rangle \in D_{i+1}$$

and, for all $a' \in \mathcal{D}$*:*

$$\langle a_1 \cdots a_i \rangle \in D_i \ \text{ if } \ \langle a_1 \cdots a_i a' \rangle \in D_{i+1}.$$

Definition 2 (Urn model) *An urn model* $\mathfrak{M}$ *is a pair* $\langle \mathcal{M}, \Delta \rangle$*, where* $\mathcal{M}$ *is a classical first-order model with domain* $\mathcal{D}$ *and* Δ *is an urn sequence over* $\mathcal{D}$*. An urn model* $\mathfrak{M}$ *satisfies a sentence* A*,* $\mathfrak{M} \vDash_u A$*, when* $\exists$ *has a winning strategy in the game* $G(A)$ *played in* $\mathfrak{M}$*.*

Definition 3 (d-invariant models) *Let* $\delta_i = \{a_i \mid \exists a_1 \cdots \exists a_{i-1} \langle a_1 \cdots a_{i-1} a_i \rangle \in D_i\}$*. For any* $d \in \mathbb{N}$*, an urn-model* $\mathfrak{M} = \langle \mathcal{M}, \langle D_1 D_2 \cdots \rangle \rangle$ *is* d*-invariant iff* $D_1 = \delta_1 = \delta_2 = \cdots = \delta_d$*.*

Here, δ_i is the set of individuals available at draw i. The d-invariant models behave as classical models for all sentences A with quantifier depth no greater than d.[18] For such A, if $\mathfrak{M} = \langle \mathcal{M}, \Delta \rangle$ is d-invariant, then $\mathfrak{M} \vDash_u A$ iff $\mathcal{M}$ classically satisfies A. Yet a valid sentence with quantifier depth d need not be satisfied by a d'-invariant model, for any $d' < d$.

Unlike relational models, changing urn models do not verify classically unsatisfiable sentences only at the cost of making some sentences both true and false (the non-subtle way!). The approach also comes with a well-defined, non-trivial notion of consequence and an accompanying proof-theory (Hintikka 1970; 1973b).[19] That's good news if one wants to investigate the logic of epistemic notions.[20] So urn models seem to be both logically respectable and good candidates for playing the role of epistemic scenarios.

There are, however, several serious problems with using urn models in this way. The first problem is this. In using these models, Hintikka works with a notion of an agent's *logical competence*, measured in terms of quantifier depth d. The rough idea is that greater competence correlates with the ability to reason correctly with sentences involving higher numbers of embedded quantifiers.[21] We could, if we want, replace Hintikka's individualistic notion of competency with a communal one, reflecting standards of linguistic understanding, or our

18. The quantifier depth of A is the greatest n such that there is a quantifier in A embedded within $n - 1$ other quantifiers (or 0 if A is quantifier-free).
19. Sequoiah-Grayson (2008) discusses Hintikka's proof theory in detail, and comes to conclusions similar to those offered below.
20. As I'll point out in §4, however, the assumption that epistemic contents are closed under logical rules (other than identity) is rather dubious.
21. In an epistemic logic setting, the epistemic accessibility relation of an agent with competency d is constrained so that all accessible urn models are d'-invariant, for all $d' \leq d$ (Hintikka 1975). That agent is then modelled as an ideal agent up until the limits of her competency, but no further.

communal expectations on rational (but non-ideal) agency.[22] But quantifier depth does not seem to be a very good measure of either logical competence or communal standards of (non-ideal) rationality. Suppose Anna completes a (correct) proof, in which no sentence has a quantifier depth greater than d, of a mathematical statement A. If she completed the proof through skill and not random luck, her achievement reflects her competence, and so we must assign her a competence of at least d. Then, no d'-invariant model was ever an epistemic possibility for Anna, for any $d' \leq d$. But since A appears in the proof, it must have a quantifier depth no greater than d, and so is true in all d-invariant models. Consequently, in a Hintikka-style model of knowledge, Anna is modelled as having known A all along; and in a model of content, her proof is modelled as being contentless and uninformative. This is just what we want to avoid.

This objection shows only that we shouldn't link the d parameter directly to an agent's competence (or communal standards of competence). Let's grant, for the sake of argument, that d can be fixed meaningfully in some other way, so as to avoid the objection. Even then, a further problem remains. Given that $\delta_1 = \mathcal{D}$ in any urn model with domain $\mathcal{D}$, we can show that, for any A which contains no embedded quantifiers and any urn model $\mathfrak{M} = \langle \mathcal{M}, \Delta \rangle$, $\mathfrak{M} \vDash_u A$ iff $\mathcal{M}$ classically satisfies A (Rantala 1975, 466, theorem 1). As a consequence, agents are modelled as being logically omniscient with respect to all such sentences; and all inferences involving only such sentences are deemed contentless. This gets things wrong: if valid deductions can be informative at all, then surely at least some purely truth-functional (quantifier-free) inferences are informative. It's not as if the first-year logic class struggles *only* with those inferences involving embedded quantification!

Hintikka's use of urn models represents an improvement (in certain respects) on using relational models to capture epistemic scenarios, but it renders too many inferences uninformative. The root of problem is that using urn models in this way establishes an absolute cut-off point between potentially informative and necessarily uninformative inferences (depending on whether they involve embedded quantifiers). But (as I argue in the next section) in reality there is no such absolute cut-off point. This is because content, and hence whether an inference is informative or not, is an inherently vague matter.

4 A Diagnosis of the Problem

In this section, I provide a diagnosis of the problem of how a valid deduction can be informative. This will pave the way for a formal model, in §5, of the content of deduction.

There is something very counterintuitive in the claim that the deductive move from, say, $A \wedge B$ to A is informative. It is tempting to say that anyone who claims to believe that $A \wedge B$ but not that A is in some way irrational (or confused about what

22. The idea, very roughly, would be that persistent logical mistakes falling below the standard would reflect the failure to grasp the relevant concept fully, whereas mistakes persistently over the threshold would instead reflect errors in calculation, lack of cognitive resources and the like.

those words mean). But if we assume (perhaps as some principle of rationality) that inferences made using such rules are wholly uninformative, we will quickly run into problems, for a deductive proof is no more than the repeated application of such rules. So if we grant the assumption, then we are at risk of incorrectly treating any proof whatsoever as being uninformative and hence incapable of adding to one's knowledge.[23] That is our puzzle.

The case is analogous to a sorites series of colour samples, going gradually from red to yellow. Adjacent samples are indistinguishable in colour and so it seems that, if we judge any sample to be red, we should also judge the next one to be red, too. But, as the first is clearly red, we are then at risk of judging them all to be red, which is clearly wrong. One could always stipulate that 'red' applies only to the first 23 samples (say), and 'non-red' to all the others. But that does nothing to resolve the puzzle, which concerns *our* concept of redness and not some artificial precisification of it. The puzzle is to make sense of truth and inference in a vague language, so that not every colour sample is counted as being red. The task is not to reform the language by removing vague predicates. I want to draw the same moral in our case of deduction: the task is to make sense of a notion of content such that some, but not all, valid inferences are informative, without drawing an artificially sharp line between those that are and those that are not. The normative notion of content we want is a vague notion, precisely because chains of seemingly uninformative inferences can give rise to informative deductions.

The analogy with sorites cases highlights another important feature of content. There must be some relation between content and meaning. In the case of a conjunction $A \wedge B$, there must be some relation between the content of $A \wedge B$, the contents of A and B, and the meaning of '$\wedge$'. Since there is a clear link between the meaning of the logical constants and the proof rules which govern their use, there must be some link between the content of $A \wedge B$ and the rules we use to make inferences to and from $A \wedge B$.[24] There is something undeniable about this: any account which wilfully ignores these rules can't claim to be a genuine account of epistemic content for rational agents.

The problematic assumption is that the way to capture such rules is in terms of closure conditions on truth-at-a-scenario. This forces us into the classical (or relational) picture of content, on which $A \wedge B$ is true at a scenario s iff both A and B are true at s. But the classical picture cannot accept that a valid inference may be informative (§1), whereas the relational models account is not acceptable for other reasons (§2). To give a non-empty content to valid deductions, we must avoid those closure conditions; yet we do not want to lose sight completely of meaning-governing inference rules. Our account of content must be subject to

23. Dummett makes a similar point: 'When we contemplate the simplest basic forms of inference, the gap between recognising the truth of the premisses and recognising that of the conclusion seems infinitesimal; but, when we contemplate the wealth and complexity of number-theoretic theorems which, by chains of such inferences, can be proved … we are struck by the difficulty of establishing them and the surprises they yield.' (1978, 297)

24. Note that I'm not claiming that proof rules are *constitutive* of a logical constant's meaning. I want to remain neutral on this issue. It may be that a logical constant's meaning is fixed by truth-conditions, which also fix the correct proof rules, for example.

those inference rules, even if truth-at-a-scenario is not closed under them. How can these requirements be satisfied simultaneously?

In the next section, I develop an account of the content of deduction on which contents are vague, rather than precise. The account avoids the closure conditions on truth-at-a-scenario which would lead to the classical (or relational model) account of content. Yet, I will claim, it gives us a notion of content that respects the inference rules which correspond to the meanings of the logical connectives.

5 A Model of Content

In this section, I set out a formal model of epistemic content (including the content of deduction). I'll begin by talking in terms of *points*, rather than scenarios. Some but not all of these points will count as epistemic scenarios, and content will be built from those scenarios. The guiding idea is that, in certain cases, it will be indeterminate whether a given point counts as a scenario. This will allow us to define contents which are themselves indeterminate.

I'll assume a space W of very fine-grained points. For each pair of arbitrary sets of sentences Γ, Δ, there is a point $w \in W$ such that the truths and falsities according to w are precisely the members of Δ and Γ, respectively. We define a model of content $\mathcal{M}$ (relative to a language $\mathcal{L}$) to be a tuple

$$\langle W, R, V^-, V^+ \rangle$$

where W is a set of fine-grained points, $R \subseteq W \times 2^W$ is an irreflexive, asymmetric and non-transitive relation and $V^-, V^+ : W \longrightarrow 2^{\mathcal{L}}$. I'll say that a point w verifies A iff $A \in V^+_w$ and falsifies A iff $A \in V^-_w$.

The relation R, which relates points to sets of points, captures the structure of proofs, relative to some fixed set of proof rules. For simplicity, I'll work with a classical sequent-style system, but this isn't essential.[25] I'll focus on the simple case of a propositional language; extension to a first-order language is easy. Standard presentations of the sequent calculus focus on *sequents*, of the form $\Gamma \vdash \Delta$, where Γ and Δ are *multisets* of sentences.[26] Rules manipulate such sequents, and may have either one or two sequents as premises (or *upper sequents*) and a single sequent as conclusions (or *lower sequents*). A standard sequent calculus for classical logic employs left and right logical rules for each connective (see, e.g., Buss 1998), plus the *contraction*, *identity* and *cut* rules:

$$\frac{\Gamma, A, A \vdash \Delta}{\Gamma, A \vdash \Delta}\ [\textsc{cl}] \qquad \frac{\Gamma \vdash A, A, \Delta}{\Gamma \vdash A, \Delta}\ [\textsc{cr}]$$

$$\frac{}{\Gamma, A \vdash A, \Delta}\ [\textsc{id}] \qquad \frac{\Gamma \vdash \Delta, A \qquad A, \Gamma \vdash \Delta}{\Gamma \vdash \Delta}\ [\textsc{cut}]$$

25. We can also encode natural deduction systems (Jago 2009c). There is no restriction to classical systems, either. We can even encode non-monotonic inference rules (Jago 2009a).
26. In a multiset, each element may occur one or more times. We can treat a multiset Γ as a standard set Δ coupled with a function $\# : \Delta \longrightarrow \mathbb{N}$, with $\#x$ telling us how many occurrences of x appear in Γ.

A proof of a sequent $\Gamma \vdash \Delta$ is a tree of sequents whose root is $\Gamma \vdash \Delta$ and whose leaves are all instances of ID. Any sequent provable using CUT can be proved without it and so I'll set CUT to one side. In practice, proofs are constructed bottom-up, beginning with the sequent to be proved and working upwards, applying the rules from lower sequent to upper sequent(s).

I'll work with a slightly non-standard system, in which (i) Γ and Δ in a sequent $\Gamma \vdash \Delta$ are sets, rather than multisets of sentences, and (ii) all sentences appearing in the lower sequent of a rule must appear in the upper sequent(s) too. The logical rules for '$\neg$' and '$\vee$', for example, are:

$$\frac{\Gamma, \neg A \vdash A, \Delta}{\Gamma, \neg A \vdash \Delta}\ [\neg\text{L}] \qquad \frac{\Gamma, A \vdash \neg A, \Delta}{\Gamma \vdash \neg A, \Delta}\ [\neg\text{R}]$$

$$\frac{\Gamma, A \vee B, A \vdash \Delta \qquad \Gamma, A \vee B, B \vdash \Delta}{\Gamma, A \vee B \vdash \Delta}\ [\vee\text{L}] \qquad \frac{\Gamma \vdash A, B, A \vee B, \Delta}{\Gamma \vdash A \vee B, \Delta}\ [\vee\text{R}]$$

The rules for the other connectives are similarly modified from the standard logical rules. In practise, these rules require us to carry all sentences with us as we construct our tree from root upwards. In this system, we can dispense with contraction (the CL and CR) rules. The system I will work with consists of just these left and right logical rules for each connective plus ID, with no structural rules. Our proof rules are thus all *strong* inference rules, in the sense of Buss (1998). This system is equivalent to the standard presentation: any sequent derivable in one system is derivable in the other. Hence the new system is sound and complete with respect to the truth-table semantics.

For every sequent $\Gamma \vdash \Delta$, there is a corresponding point w which verifies each sentence of Γ and falsifies each sentence of Δ. We can, therefore, write that sequent as $V_w^+ \vdash V_w^-$. This sequent is valid iff w is logically incoherent, for $\Gamma \vdash \Delta$ iff the combined truth of each $A \in \Gamma$ is incompatible with the combined falsity of each $B \in \Delta$. R is then the set of all pairs $(w, \{u\})$ for which there are points w and u and a rule instance r_1:

$$\frac{V_u^+ \vdash V_u^-}{V_w^+ \vdash V_w^-}$$

plus all pairs $(w, \{u, v\})$ for which there are points w, u and v and a rule instance r_2:

$$\frac{V_u^+ \vdash V_u^- \qquad V_v^+ \vdash V_v^-}{V_w^+ \vdash V_w^-}$$

In this way, R relates points corresponding to a rule instance's lower sequent to the set of points corresponding to its upper sequent(s). Each R-transition corresponds to a potential inference, all of which involve logical (as opposed to structural) rules. R as a whole captures all the potential proofs available to us.

Let a *point-graph* G be any rooted, directed acyclic graph with vertices $V_G \subseteq W$ and edges $E_G \subseteq W^2$, restricted so that $(w, u) \in E_G$ only if:

$$\exists X \subseteq W((w, X) \in R \ \&\ u \in X \ \&\ \forall x(x \in X \leftrightarrow (w, x) \in E_G)).$$

In a point-graph G, each vertex has zero, one or two children. Vertices with zero children are *leaves*. If w's only child is u, then there is an instance of a proof rule such as r_1 above. If w has two children u, and v, then there is an instance of a proof rule such as r_2 above. A point w is *closed* iff $V_w^+ \cap V_w^- \neq \varnothing$, and a *point-proof* is a point-graph all of whose leaf nodes are closed. The *size* of a point-graph G is the number of non-leaf vertices it contains (which corresponds to the number of inference tokens G captures). Each closed point w is an instance of the ID rule, and each point-proof has an instance of ID at each of its leaves. Given the soundness of our rules and the way point-proofs are constructed from rule instances, it follows that for any point-proof P, the point w at the root of P is logically inconsistent (and so the sequent $V_w^+ \vdash V_w^-$ is valid). Conversely, for any point w, if w is logically inconsistent then, given the completeness of the proof rules and the way point-proofs are constructed by rule instances, there is a point-proof P with w at its root.

Next, we totally order all our points. Let fw be the size of the smallest point-proof with w at its root, if there is such a point-proof (and undefined otherwise). Now set $w \leq w'$ iff either fw' is undefined or $fw \leq fw'$. Points w such that $V_w^- \cap V_w^+ \neq \varnothing$ are $\leq$-minimal elements, associated with point-proofs of size 0; I'll call these the $\perp$-*points*. Points not at the root of any point-proof are $\leq$-maximal. The maximal points are consistent (with respect to our chosen proof rules). All other points are inconsistent; but not all inconsistent points are on a par. Intuitively, $\leq$ orders points by how easy (in terms of proof length) it is to refute that point. When $w \leq w'$, the inconsistencies in w' are buried at least as deep (in terms of proof length) as those in w. So, if one accepts w as an epistemic scenario, then one should also accept w' as an epistemic scenario.

This does not yet give us an account of which points count as epistemic scenarios, which we need for our account of content. The ordering $\leq$ is supposed to correspond to the way we can order colour samples by their degree of redness. The problem in that case is that we can't detect just which set of samples count as the extension of 'red'. The colour case and the deduction case are structurally similar, and so I want to treat them in the same way. Whatever one thinks is the correct philosophical account of vagueness can be 'plugged in' at this point.

I'm tempted by the view that vague cases are cases for which there is no fact of the matter either way. On this view, we partition the points into three classes: those that are scenarios, those that are not, and those for which there is no fact of the matter either way. This partition is constrained by $\leq$: if w is a scenario, then so are all w' such that $w \leq w'$; if w is not a scenario, then so are all w' such that $w' \leq w$; and if there are no facts of the matter regarding w_1 and w_2, and $w_1 \leq w_2$, then there is no fact of the matter regarding any w' such that $w_1 \leq w' \leq w_2$. One will reason about this set-up using a 3-valued logic, e.g. strong Kleene logic. Epistemicists, by contrast, will partition points into two: the scenarios and the non-scenarios, again constrained by $\leq$ in the obvious way. Many-valued accounts will assign a degree of truth $\delta_w \in [0,1]$ to 'w is a scenario', constrained so that $\delta_w \leq \delta_{w'}$ iff $w \leq w'$.[27]

27. For an account along these lines, see Jago 2009b.

Epistemic contents inherit the vagueness of 'epistemic scenario'. Let $|A|^+$ be the set of all epistemic scenarios which verify A, and $|A|^-$ be the set of all epistemic scenarios which falsify A.[28] These sets have indeterminate membership: $s \in |A|^+$ iff (i) s verifies A and (ii) s is a scenario. Since it may be indeterminate whether s is a scenario, it may also be indeterminate whether $s \in |A|^+$ (and similarly for $|A|^-$). Just how this vagueness in content is modelled depends, once again, on one's semantics for vagueness.[29] The content of A is then the pair $(|A|^+, |A|^-)$. Content thus tracks all scenarios which falsify, as well as all those which verify, the sentence in question. Finally, I define the content of the deduction from Γ to A as the set of all epistemic scenarios which verify each of Γ and falsify A, i.e. $|\Gamma|^+ \cap |A|^-$.[30]

In this section, I've presented models which can be used to define epistemic contents of sentences and of deductions. Those contents are vague: it may be indeterminate whether a content has a given point w as a member. Just how this vagueness is modelled depends on one's semantics for vagueness in general: the models given here can be supplemented by any of the formal semantics for vague languages. In the next section, I discuss the features and advantages of this account of content.

6 Evaluating the Model

In this section, I review some of the advantageous features of the account I've just presented. To begin with, it improves on the urn-models account (§3) in that it allows purely truth-functional deductions to count as informative. On the view I've presented, whether a given deduction is informative isn't a matter of whether it is contained within this or that fragment of the language; rather, it is a matter of how difficult that inference is.[31] It also improves on the relational models approach (§2) in that no obviously impossible point is counted as an epistemic scenario, and so no such point features in the content of any deduction.

On the account I've given, truly trivial deductions come out contentless and hence uninformative, just as we want. For example, suppose w verifies each of

28. Note that, in general, $|A|^- \neq |\neg A|^+$.

29. For epistemicists, $|A|^+$ and $|A|^-$ are classical sets whose exact membership (when so described) is unknowable. On the many-valued approach, $|A|^+$ and $|A|^-$ are fuzzy sets, whose membership function is δ from above.

30. This is a variant on content$_2$, defined as $|\Gamma|^+ \cap |\neg A|^+$. Defining content as I have done makes sense even in system where negation behaves in non-standard ways (e.g., in paraconsistent logic). Working with content$_1$ has bad consequences. Take $A, B \vdash A \wedge B$: we find its content$_1$ simply by finding scenarios which verify both A and B but say nothing about $A \wedge B$. There exist such scenarios so long as A and B are mutually consistent (for example, the incomplete but consistent point which verifies A and B but nothing else counts as a scenario). This is a rather badly-behaved notion of content. For example, $p, q \vdash p \wedge q$ counts as contentful$_1$, whereas $p, \neg p \vdash p \wedge \neg p$ does not (because any point verifying p and $\neg p$ does not count as a scenario), despite each inference being an instance of *conjunction introduction*. Similarly, $A \vdash A$ is deemed contentless$_1$ (as it should be), whereas $A \vdash A \wedge A$ and $A \vdash A \vee A$ are not. This strikes me as a bizarre position to hold on content.

31. More precisely: it is a factor of the shortest number of inference steps required to move from premises to conclusion, relative to some fixed set of inference rules.

$A \to B$, A and falsifies B. Then there is a point-proof with w at its root, which we can represent as

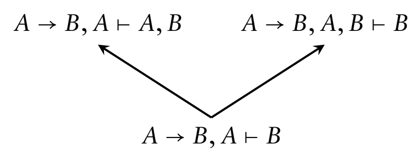

This graph has size 1, hence $fw \leq 1$. No such point counts as an epistemic scenario, hence the content of the deduction from $A \to B$ and A to B, defined as $|A \to B|^+ \cap |A|^+ \cap |B|^-$, is empty.

Yet not all valid deductions are empty. Suppose that points w for which $fw \geq m$ count as epistemic scenarios. Then the deduction

$$p_1, p_1 \to p_2, p_2 \to p_3, \ldots, p_{n-1} \to p_n \vdash p_n$$

is contentful when $n > m$. Its content consists of scenarios verifying p_1 and each $p_i \to p_{i+1}$ $(i < n)$ and falsifying p_n. We can verify that there exist such scenarios as follows. Let w be such that $V_w^+ = (\bigcup_{i<n}\{p_i \to p_{i+1}\}) \cup \{p_1\}$ and $V_w^- = \{p_n\}$. The shortest point-proof with w at its root corresponds to $n - 1$ applications of $\to$L, hence $fw = n - 1 \geq m$ and so, by assumption, w counts as a scenario. There are then infinitely many scenarios verifying p_1 and each $p_i \to p_{i+1}$ $(i < n)$ and falsifying p_n: to construct one, simply extend V_w^+ or V_w^- (or both) in a way which does not allow for a smaller point-proof to be constructed than the one just considered.

We have an account of content on which some, but not all, valid inferences are informative. What of the additional requirement, discussed in §4, that the content of a logically complex sentence should in some way be linked to the proof rules governing the use of its main connective? Let's focus, for the moment, on the positive component $|A|^+$ of sentence A's content. In a classical possible-worlds system of content, we would have that

(1) $|A \wedge B|^+ \subseteq |A|^+$ and $|A \to B|^+ \cap |A|^+ \subseteq |B|^+$.

These inclusion relationships capture (one aspect of) the meaning of '∧' and '→'. But these inclusion relationships do not hold in our present epistemic system, for epistemic scenarios are not closed under *conjunction elimination* or *modus ponens*. We can find scenarios w which are members of $|A \wedge B|^+$ but not $|A|^+$. Nevertheless, our epistemic notion of content does capture an aspect of these classical inclusion relationships. The classical possible-worlds framework identifies $(|A|^-)^c$ (the set-theoretic complement of $|A|^-$) with $|A|^+$ and so, on a domain of classical possible worlds, $|A|^+ \subseteq |B|^+$ holds iff $|A|^+ \subseteq (|B|)^c$ holds. Thus on a domain of classical possible worlds, the inclusion relationships in (1) are equivalent to

(2) $|A \wedge B|^+ \subseteq (|A|^-)^c$ and $|A \to B|^+ \cap |A|^+ \subseteq (|B|^-)^c$.

Although epistemic space does not verify the inclusion relationships in (1), it does verify those in (2).[32] This is one way in which this epistemic notion of content captures the meanings of '$\wedge$' and '$\rightarrow$'. Similar things can be said for the other connectives.

In fact, although epistemic space does not validate (1), it nevertheless enforces a tight relationship between $|A \wedge B|^+$ and $|A|^+$ and between $|A \rightarrow B|^+ \cap |A|^+$ and $|B|^+$ (and similarly for the other connectives). Let's focus on the region r of $|A \wedge B|^+$ not included in $|A|^+$. Some of the scenarios in that region might be just one inference away from non-scenarios (or from indeterminate cases of scenarios). Set such scenarios to one side. For all the remaining scenarios w, there is a further scenario $u \in |A|^+$ such that $(w, \{u\}) \in R$. Intuitively, such scenarios are as close as they could be to $|A|^+$, without actually being in that region. In this sense, $|A \wedge B|^+$ comes as close as it could be to $|A|^+$ without being included in $|A|^+$. We can say similar things about the relationships between $|A \rightarrow B|^+ \cap |A|^+$ and $|B|^+$, between $|A|^+$ and $|A \vee B|^+$, and so on.[33] This 'closeness' relationship between the relevant contents justifies the claim that those contents respect the relevant inference rules. And since those inference rules are intimately connected to the meanings of the logical constants, this in turn justifies the claim that this model of content respects the meanings of the logical constants.

In summary, we have a model of content which counts some, but not all, valid deductions as informative. In particular, completely trivial deductions are modelled as contentless and hence uninformative. Content (and hence informativeness) is treated as a vague notion: it is indeterminate just which points constitute a particular content. Moreover, our model preserves the intimate relation between the content of $A \wedge B$ and the contents of A and B (and similarly for other connectives). Those contents are as close as they could be, without collapsing into the classical picture of content, which is unable to model informative inference.

References

Buss, Samuel. (1998). *Handbook of proof theory*, volume 137, Elsevier, Amsterdam.

Chalmers, David. (2002). 'The components of content', *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. Chalmers, Oxford University Press, pp. 608–633.

Chalmers, David. (2010). 'The nature of epistemic space', *Epistemic Modality*, ed. A. Egan and B. Weatherson, Oxford University Press.

Dummett, Michael. (1978). 'The justification of deduction', *Truth and other enigmas*, Harvard University Press, Cambridge, MA, pp. 166–185.

32. Of course, it is not the case that $|A|^+ \subseteq (|B|^-)^c$ holds whenever $|A|^+ \subseteq |B|^+$ holds on a classical possible-words domain. For some hard-to-prove tautology $\top$, for example, we can have $|A|^+ \nsubseteq (|\top|^-)^c$.
33. We cannot say something similar about $|A|^+$ and $|B|^+$ whenever $A \vdash B$. For if the shortest proof from A to B is long, then some scenarios in $|A|^+$ might be quite a long way off (in terms of R-transitions) from any scenario $|B|^+$. But this is the point of the epistemic notion of content: as proofs from A to B become harder to spot, the content of A becomes more remote from the content of B. This is precisely why such inferences are informative, whereas inferences such as $A \wedge B \vdash A$ are not.

Henkin, Leon. (1961). 'Some remarks on infinitely long formulas', *Infinitistic Methods*, Pergamon Press, Oxford, pp. 167–183.

Hintikka, Jaakko. (1962). *Knowledge and belief: an introduction to the logic of the two notions*, Cornell University Press, Ithaca, N.Y.

Hintikka, Jaakko. (1970). 'Surface information and depth information', *Information and Inference*, ed. J. Hintikka and P. Suppes, Reidel, Dordrecht.

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

Hintikka, Jaakko. (1973b). 'Surface semantics and its motivation', *Truth, Syntax and Modality*, ed. H. Leblanc, North-Holland, Amsterdam.

Hintikka, Jaakko. (1975). 'Impossible possible worlds vindicated', *Journal of Philosophical Logic* 4: 475–484.

Jago, Mark. (2009). 'Epistemic logic for rule-based agents', *Journal of Logic, Language and Information*, 18(1):131–158.

Jago, Mark. (2009). 'Logical information and epistemic space', *Synthese*, 167(2):327–341.

Jago, Mark. (2009). 'Resources in epistemic logic', in J.-Y. Béziau and A. Costa-Leite, editors, *Dimensions of Logical Concepts*, volume 55, pages 11–33. Coleção CLE, Campinas, Brazil.

Lakemeyer, Gerhard. (1986). 'Steps towards a first-order logic of explicit and implicit belief', *Proceedings of the First Conference on Theoretical Aspects of Reasoning About Knowledge*, ed. J. Y. Halpern, Morgan Kaufmann, San Francisco, pp. 325–340.

Lakemeyer, Gerhard. (1987). 'Tractable metareasoning in propositional logic of belief', *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pp. 401–408.

Lakemeyer, Gerhard. (1990). 'A computationally attractive first-order logic of belief', *Proceedings of JELIA 90*, Springer, Heidelberg, pp. 333–347.

Levesque, Hector. (1984). 'A logic of implicit and explicit belief', *Proceedings of the Fourth National Conference on Artificial Intelligence*, pp. 198–202.

Lewis, David. (1975). 'Language and languages', *Language, Mind and Knowledge*, ed. K. Gunderson, University of Minnesota Press, pp. 3–35.

Lewis, David. (1986). *On the Plurality of Worlds*, Blackwell, Oxford.

Peirce, Charles. (1992). *Reasoning and the Logic of Things: The Cambridge Conferences Lectures of 1898*, Harvard University Press, Cambridge Mass.

Priest, Graham. (1987). *In Contradiction: A Study of the Transconsistent*, Martinus Nijhoff, Dordrecht.

Priest, Graham. (2008). *An Introduction to Non-Classical Logic*, Cambridge University Press, Cambridge.

Rantala, Veikko. (1975). 'Urn models', *Journal of Philosophical Logic* 4: 455–474.

Sequoiah-Grayson, Sebastian. (2008). 'The scandal of deduction', *Journal of Philosophical Logic* 37(1): 67–94.

Stalnaker, R. (1976). 'Propositions', *Issues in the Philosophy of Language*, ed. A. MacKay and D. Merrill, Yale University Press, New Haven, pp. 79–91.

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

van Benthem, J. (2011). *Logical dynamics of information and interaction*, Cambridge University Press, Cambridge.

van Benthem, J. and Martinez, M. (2008). 'The stories of logic and information', *Handbook of the Philosophy of Information*, ed. J. van Benthem and P. Adriaans, Elsevier, Amsterdam, pp. 217–280.