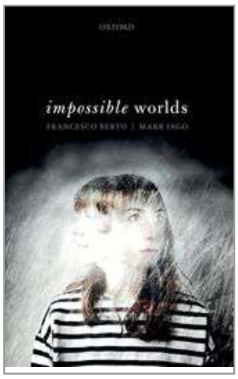

Impossible Worlds
Francesco Berto and Mark Jago

Print publication date: 2019
Print ISBN-13: 9780198812791
Published to Oxford Scholarship Online: August 2019
DOI: 10.1093/oso/9780198812791.001.0001

Information and Content

Francesco Berto
Mark Jago

DOI:10.1093/oso/9780198812791.003.0009

Abstract and Keywords

This chapter conceptualizes information in terms of ruling out scenarios. It discusses informative identity statements, which give rise to Frege’s puzzle, and the problem understanding how a valid logical inference can be informative. An analysis of informative logical inferences is given, on which the content of a valid deduction is often indeterminate. A consequence is that it is indeterminate exactly which logical inferences are informative. The chapter then analyses a rather different notion of content, concerning *what is said* by a speaker in making an utterance.

Keywords: informative identity statements, Frege’s puzzle, informative inference, what is said

9.1 Informative Statements

Here are some informative statements. It’s currently cold but dry in Nottingham. James Newell Osterberg is Iggy Pop. Fermat’s Last Theorem is true. If all truths are knowable, then all truths are known. And here, by contrast, are some uninformative statements. No bachelor is married. Iggy Pop is Iggy Pop. 1 + 1 = 2. Either all truths are knowable or they’re not.

The first list of statements might be informative to some people and not to others. For someone standing outside in Nottingham right now, it’s probably not informative to be told that it’s currently cold but dry there. For someone who knows lots about Iggy Pop, it’s probably not informative to be told that James Newell Osterberg is Iggy Pop. To those who’ve encountered Fitch’s paradox, it’s probably not informative to be told that, if all truths are knowable, then all truths are known. To modern mathematicians, it's probably not informative to be told that Fermat's Last Theorem is true.

Whether a statement is informative to someone depends on what information that person already has. It also depends on the way in which they have that information. Take our person standing outside in Nottingham right now. She may be confused about where she is. If she doesn't know she's in Nottingham, then experiencing the weather in her immediate surroundings won't help her to conclude that it's currently cold and dry in Nottingham. Then her weather app, which says *currently cold and dry in Nottingham*, might be informative to **(p. 186)** her. In some sense, one might think, she already knew this. She knew it's currently cold and dry where she is, and Nottingham is where she is. But still, it was informative to her to be told that it's currently cold and dry in Nottingham. What's informative depends also on how the information is presented.

In this chapter, we investigate accounts of what it is for a statement to be informative. We're interested in what information is, in and of itself. We're also interested in how a statement gets to be informative. As the examples above suggest, logical and mathematical truths can be informative (to some people, at least). This in turn hints that an impossible worlds framework is a promising way to understand this phenomenon. Yet, as the examples also suggest, not all logical and mathematical truths are informative (to anyone). So, if we go down the impossible worlds route, we need to be careful about *which* impossible words are included in the analysis.

Similarly, some but not all identity statements are informative (to some people). If identity is necessary, as most philosophers hold after Kripke (1980), then impossible worlds seem to be an attractive way to go. But again, we will need to be careful about which impossible worlds feature in the analysis. Informativeness, it will turn out, is a very puzzling concept.

9.2 Information as Ruling Out Scenarios

According to a popular analysis, for a sentence (or the expressed proposition) to be informative is for it to rule out certain scenarios, or would-be possibilities. Hintikka (1962) gave the classic presentation of this view, which Chalmers (2002a, 2010), Lewis (1975, 1986b), and Stalnaker (1976b, 1984) then put to work in various ways. Van Benthem (2011) and Van Benthem and Martinez (2008) discuss the recent literature.

The proposition *that it often rains in Manchester* is informative because it excludes scenarios in which it rains infrequently in Manchester. Before an agent comes to believe that proposition, it was **(p.187)** possible, as far as she was concerned, that it rains infrequently in Manchester. In coming to believe that proposition, she ceases to treat such scenarios as ways the world might be, for all she knows.

We can think of all scenarios according to which it often rains in Manchester as constituting a notion of content for ‘it often rains in Manchester’ which is suitable for various epistemic purposes. To believe that proposition is to treat *only* those scenarios according to which it often rains in Manchester as being doxastic possibilities; and to know that proposition is to treat *only* those scenarios according to which it often rains in Manchester as being epistemic possibilities (Chapter 5). 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.

If the scenarios in question are all possible worlds, then problems ensue. In §1.2, we introduced the Bar-Hillel-Carnap theory of information (Bar-Hillel and Carnap 1953). This claims that the informative role of a sentence consists in splitting the totality of possible worlds into those where it is true, and those where it is false. The consequence outlined in §1.3 is that identity statements of the form $a = b$, and metaphysical, logical, and mathematical truths, all end up being treated as uninformative.

As a consequence, the Bar-Hillel-Carnap theory denies that any logical deduction or mathematical proof can ever be informative. This is implausible as a characterization of information for limited, resource-bound, and fallible cognitive agents like us all. The issue, which Floridi (2015, §4.1) calls the ‘Bar Hillel-Carnap paradox’, is in fact just a variation on the logical omniscience problem in epistemic logic (Chapter 5), which we’ll discuss further in Chapter 10.

Acknowledging the problem, Bar-Hillel and Carnap say that their account should not

> be understood as implying that there is no good sense of ‘amount of information’ in which the amount of information of these sentences will not be zero at all, and for some people, might even be rather high.
>
> (Bar-Hillel and Carnap 1953, 229)

(p.188) This ‘good sense’ is a kind of ‘psychological information’ (1953, 229, on which the Bar-Hillel-Carnap theory has nothing further to say. On this view, there is one kind of information, semantic (and, presumably, non-psychological) information, which applies to contingent, empirical statements. And there is distinct ‘psychological’ notion which, in some unspecified way, makes sense of the informativeness of certain necessary truths.

That approach looks messy and poorly motivated. We don’t find pre-theoretical reasons for wanting to divide notions of information along these lines. Learning about the logical consequences of some supposition, or of a putative move in a game of chess, seems to us to be informative, in much the same sense that contingent statements can be informative. If asked whether it’s currently raining in Sydney, or whether Qe4 is a good next move, an uninformed agent might in each case say ‘could be’. In each case, becoming informed rules out would-be possibilities for the agent in question. Both cases link to belief and knowledge in the same way. In each, gaining information leads to new beliefs and, in the right circumstances, to fresh knowledge. Moreover, both cases have a psychological element. Both can be surprising; both contents can be the objects of an agent’s hopes or fears; both may interact with the agent’s emotions. In short, we think a unified notion of semantic informational content, one which can deal with both cases, is preferable.

We don’t thereby want to claim that there is only one good notion of information. Far from it: there are many such notions, and more than one may be theoretically useful. In §9.1, we contrasted the concept of being *potentially informative to some agent* with *being informative to agent x*, for a specified *x*. In §9.6, we’ll discuss a distinct notion of informational content, which concerns *what is said* by a speaker in making an utterance. Our overall strategy is the one we discussed in §8.4. We start with fine-grained worlds. We impose additional *inter*-world structure, or *intra*-world closure conditions, depending on the concept of information under investigation.

(p.189) 9.3 Informative Identities

James Newell Osterberg recently turned 70. That information would be of little interest to us (to Franz and Mark, at least) if we didn’t also know that James Newell Osterberg is Iggy Pop. Given what we know about his time in The Stooges, it’s somewhat surprising that he’s made it to 70. If we didn’t know that James Newell Osterberg is Iggy Pop, we’d not find it so surprising that James Newell Osterberg has reached 70. ‘James Newell Osterberg is Iggy Pop’ is an *informative identity*. It allows us to connect our attitudes to particular bits of information to other bits of information. It allows us to connect our *Iggy Pop* information, and the attitudes we take to it, to our *James Newell Osterberg* information.

How can ‘James Newell Osterberg is Iggy Pop’ be informative? Since it’s true, ‘Iggy Pop’ is a name for Osterberg, as is ‘James Newell Osterberg’. Both names pick out the same individual. So, semantically, ‘James Newell Osterberg is Iggy Pop’ would seem to express the proposition *that Osterberg is Osterberg*. But that’s utterly trivial and uninformative. So how are we to understand the information conveyed by ‘James Newell Osterberg is Iggy Pop’?

This is a version of *Frege’s puzzle* (Frege 1892). Frege’s own solution was (in effect) that propositions are not identified by the worldly entities they are about, but rather by the *modes of presentation* of those worldly entities (he called them *Sinne*, senses). If the modes of presentation of ‘*a*’ and ‘*b*’ differ, then the Fregean propositions (which Frege called 'thoughts') *that Fa* and *that Fb* differ. This allows an agent to stand in a relation (such as *believing* or *knowing*) to one proposition but not the other. So, in particular, one can learn that Iggy Pop has turned 70, even if one already knew that Osterberg has turned 70. In this way, Frege can explain how replacing 'Osterberg' with 'Iggy Pop' can be informative. In particular, he can explain how 'Osterberg is Iggy Pop' is informative, whereas neither 'Osterberg is Osterberg' nor 'Iggy Pop is Iggy Pop' are.

To support this approach, Frege requires some rather elaborate semantic machinery. On a straightforward semantic account, '*a*' refers **(p.190)** to *a*, '*b*' to *b*, and it is these referents, *a* and *b*, which feature in the truth-conditions for the utterance in question. And so it goes on Frege's analysis, in *direct* discourse. But in *indirect* discourse, such as the sentence following 'it is informative that' or 'believes that', '*a*' and '*b*' refer not to *a* and *b* but to their mode of presentation. It is these modes of presentation which feature in the truth-conditions of the utterance as a whole. Thus, on the Fregean view, indirect contexts induce a switch of reference, from the usual worldly entities to their modes of presentation.

The Fregean approach is rich and powerful. It's not our intention here to evaluate it. Our topic is impossible worlds. Our aim is to show that a worlds-based approach can analyse informative identities (and belief reports: see Chapter 10) as well as the Fregean approach.

If we're successful, then we can avoid a question that troubles the Fregean approach: just what are senses? Frege speaks of the 'realm of senses', distinguished from the 'realm of reference' (Frege 1956). It seems that, for Frege, senses are *sui generis* entities, neither physical nor mental (see Dummett 1993a, 154). Dummett notes that, for Frege, 'the realm of sense is a very special region of reality; its denizens are, so to speak, things of a very special sort' (Dummett 1993a, 154). Senses must be entities of some kind, else they could not be referents in belief-contexts (as Frege's theory claims). But if they are primitive, non-causal abstract entities, how can we refer to them in belief reports?

There may well be good answers to these questions. (Chalmers (2002c) identifies Fregean senses with 'primary' or 'epistemic' intensions, a particular function from possibilities to extensions. But Chalmers's approach makes no concession to hyperintensional notions.) We avoid the worry entirely if we can show that a worlds-based approach is up to the job. (Indeed, Bjerring and Rasmussen (2017) and Jago (2014a) suggest that the best approach to understanding senses is in terms of functions on possible and impossible worlds.)

Our alternative to the Fregean approach goes as follows. If we accept Nolan's Principle (NP), or any of the other principles from §8.4, then for any non-empty terms 'a' and 'b', there are guaranteed **(p.191)** to be worlds which represent that $a = b$ and worlds which represent that $a \neq b$. For suppose $a = b$. Then the actual world represents that $a = b$. And since it is impossible that $a \neq b$, by (NP) some impossible world represents that $a \neq b$. If we suppose instead that $a \neq b$, then the argument is similar. (Just *how* a world represents that $a \neq b$, given that in fact $a = b$, is a further question. Jago 2014a, §§5.5–5.6 is one attempt at a solution.)

Since there are guaranteed to be worlds according to which $a \neq b$, a true identity statement '$a = b$' is guaranteed to have a non-empty content. An agent may take any doxastic attitude to that content: belief, disbelief, or neither. If she does not believe it, it is because she takes some of the worlds where a is not b (which are, in fact, impossible) to be ways the world could be, for all she knows. This is compatible with her being a rational agent and perfectly competent language-user (she might even be a heavily idealized agent who knows all a priori truths). So there is no rational compulsion for her to believe that content. She may go from not believing it to believing it and, in the right circumstances, she may gain it as knowledge in the process. If she does, it is informative to her.

On this approach, true identity statements '$a = b$' are potentially informative. This approach maintains the benefits of the Fregean approach, but without relying on reference switching mechanisms. (In Chapter 10, we'll also argue that the impossible worlds approach can make sense of belief ascriptions without resorting to reference switching.)

Won't the approach incorrectly treat '$a = a$' as being informative, too? The worry arises because $a = a$ is a logical truth, and hence it's logically impossible that $a \neq a$. So, given (NP) or one of the stronger principles from §8.4, there are worlds which represent that $a \neq a$. Mustn't we then treat '$a \neq a$' and being potentially informative?

We can resist this final move. Some contents are not suitable objects of epistemic attitudes. Not all impossible worlds are epistemically possible: some are not epistemically accessible for any agent. Some such worlds represent blatant contradictions, like representing some A as being both true and false. On the account of logical information **(p.192)** we offer in §9.5, such worlds are *deeply* epistemically impossible. No such world is accessible to any possible agent. The details of the view will have to wait until §9.5. But we can already anticipate that, in just the same way, worlds which represent some a as not being self-identical are deeply epistemically impossible. So, although there are sets of worlds which represent that $a \neq a$, no such set of worlds is an *epistemic* content.

To be informative, a statement 'A' must be capable of being disbelieved. It must be possible for an agent to believe that A, or to believe that $\neg A$ instead. But since $a \neq a$ cannot be believed (on the account we're suggesting), it follows that $a = a$ is treated correctly as an uninformative identity.

9.4 Informative Inference

This section draws on Jago (2013b). Deductive reasoning is essential to philosophy, mathematics, and logic. In those areas and others, its use is beyond question, and this must be so, at least in part, because of information it conveys. But *how* can deduction carry information, if, in some sense, the premises already guarantee the conclusion? In 'The Justification of Deduction', Dummett (1978a, 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.

How, wonders Dummett, can the move from premises to conclusion be informative, if the former already guarantee the latter? It is 'a delicate matter so to describe the connection between premisses and conclusion as to display clearly the way in which both requirements [justification and usefulness] are fulfilled' (Dummett 1978a, 297). 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.

We might think of the information content of a valid deduction $\Gamma \vdash A$, from premises Γ to conclusion A, in terms of the differences **(p.193)** an agent's belief state might undergo in performing that deduction. We can consider an agent who initially believes the premises but not the conclusion, and who ends up believing the conclusion (on the basis of the deduction she performs). Alternatively, we can think in terms of an agent discovering the incompatibility of the premises Γ with the conclusion's falsity. Either way, we are analysing some relationship between the content of the premises and the content of the conclusion.

Let's use the notion '$|A|$' to denote the set of worlds (possible or impossible) which represent that A. For sets of sentences Γ, we'll use '$|\Gamma|$' to denote the set of worlds which represent that B, for each $B \in \Gamma$. Our approach throughout this book has been to analyse notions of content in terms of possible and impossible worlds. In our present setting, the worlds in question have to outrun the logically possible ones. For suppose we limit each set $|A|$ to the possible worlds. Then if Γ entails A, $|\Gamma|$ already includes $|A|$, and so already excludes A's being false. So, however we analyse the relationship between premise and conclusion contents, we will be working with sets that include logically impossible worlds. As a minimal requirement, what they represent must not be closed under classical logical consequence.

A popular place to look for such worlds is the model theory of *paraconsistent* and *paracomplete* logics, which we encountered in §5.4 in the guise of *FDE worlds*. At an FDE world, a sentence may be true, false, both, or neither. This is accomplished by replacing the usual valuation function with a relation, ρ, which may relate a sentence to 1 or 0, to both, or to neither.

Let's take the content of premises and conclusion to be given in terms of such worlds. Our first notion of content of an inference $\Gamma \vdash A$ focuses on the difference between the premises without the conclusion and the premises with the conclusion. This amounts to those worlds according to which the premises are true, but the conclusion is not: $|\Gamma| - |A|$. Call this *content*$_1$. Our second notion analyses the content of $\Gamma \vdash A$ in terms of those worlds where the premises are true but the conclusion is false. In FDE worlds, A is false iff $\neg A$ is true, and so this notion of content amounts to $|\Gamma| \cap |\neg A|$. Call this notion *content*$_2$.

(p.194) These notions are classically equivalent but differ in our paraconsistent and paracomplete FDE setting, since, as we saw in §5.4, being false and failing to be true come apart at FDE worlds. We shall say that an inference is *trivial*$_1$ (or *trivial*$_2$) just in case its content$_1$ (or content$_2$) is the empty set. We'll use 'non-trivial$_{1/2}$' and 'contentful$_{1/2}$' interchangeably, and we'll reserve 'trivial', without a subscript, to capture the non-technical sense in which inferences like $A \vdash A$ (but not all valid inferences) seem obvious and uninformative.

FDE models are not in general closed under *modus ponens* for the material conditional $\supset$: there are worlds w where both $A \supset B$ and A are true, but B is not: $\rho_w(A \supset B)1$ and $\rho_w A1$ but not $\rho_w B1$. So $|\{A \supset B, A\}| - |B|$ is non-empty: *modus ponens* on $\supset$ is non-trivial$_1$. Similarly, it is non-trivial$_2$, since there are worlds w where both $A \supset B$ and A are true, but B is false: $\rho_w(A \supset B)1$, $\rho_w A1$ and $\rho_w B0$.

On this picture, not all valid deductions come out as being contentful$_1$. The deduction $A, B \vdash A \wedge B$ remains trivial$_1$, since any FDE world verifying A and B individually also verifies $A \wedge B$, and so $|\{A, B\}| - |A \wedge B|$ is empty. Indeed, any classically valid inference whose only connectives are '$\wedge$' and '$\vee$' will be deemed trivial$_1$, on this view. This is a puzzling feature 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, then why not the other?

By contrast, *every* valid inference is deemed contentful$_2$. In FDE worlds, the conclusion may be both true and false. So even where the premises guarantee the truth of the conclusion in our FDE setting, they do not thereby rule out its falsity. Even the most seemingly trivial inference of all, $A \vdash A$, is deemed contentful$_2$. Its content$_2$ is $|A| \cap |\neg A|$, which is the set of all worlds w where A is both true and false: $\rho_w A1$ and $\rho_w A0$. This is an even worse consequence than the results for content$_1$. Surely some inferences are so trivial as to contain no information whatsoever. If $A \vdash A$ is deemed informative, then we seem to have a worthless notion of information.

There is a deeper problem with the FDE worlds approach: it fails to explain why the worlds it provides are suitable tools for analysing **(p.195)** 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 worlds, which assign both 0 and 1 to some sentence. To see why, assume that $\Gamma \vdash A$. Then for any consistent assignment ρ_w on which $\rho_w B1$ for each $B \in \Gamma$, $\rho_w A1$ too, and hence (given consistency) not $\rho_w \neg A1$. But then $w \notin (|\Gamma| - |A|)$ and $w \notin (|\Gamma| \cap |\neg A|)$. So each notion of content can contain only explicitly contradictory worlds, at which some A is both true and false.

The problem is that it is hard to see why such explicitly contradictory worlds should play a role in an epistemic notion of content. If what a world represents is obviously impossible to any agent who meets minimal standards of rationality, then there's no sense in which ruling out that world corresponds to gaining new information.

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 worlds, then we have to admit impossible worlds. But obviously impossible worlds, representing explicit contradictions, cannot feature in any account of rational attitudes. And on the FDE-worlds account of deductive content, the obviously impossible worlds are *all* we're left with. In short, our question is difficult because it requires us to find worlds which are impossible, but not obviously so.

Our problem, therefore, is not merely to find worlds not closed under classical consequence. The problem is to provide a notion of a world which is logically impossible, but not obviously so. Lewis (in arguing against paraconsistent logic) puts the point nicely:

> I'm increasingly convinced that I can and do reason about impossible situations. … But I don't really understand how that works. Paraconsistent logic … allows (a limited amount of) reasoning about *blatantly* impossible situations. Whereas what I find myself doing is reasoning about *subtly* impossible situations, and rejecting suppositions that lead fairly to blatant impossibilities.
>
> (Lewis 2004, 176)

On Lewis's analysis, 'make-believedly possible impossibilities' might well have a use in the analysis of content, but:

> **(p.196)** The trouble is that all these uses seem to require a distinction between the subtle ones and the blatant ones (very likely context-

dependent, very likely a matter of degree) and that's just what I don't understand.

(Lewis 2004, 177)

Hintikka (1975), whilst addressing the logical omniscience problem head-on, makes a similar point (see §5.3). He argues that, for epistemic purposes, impossible worlds must be 'subtly inconsistent' worlds which 'look possible but which contain hidden contradictions' (Hintikka 1975, 476–8). The core problem with FDE worlds (and will all similar approaches) is that they are either logically possible, or blatantly impossible.

How can we make sense of a world being subtly impossible? We present one attempt in the next section.

9.5 Vague Logical Information

Jago (2013b, 2014a) argues that we should view the problem of informative inference as an instance of the problem of vagueness. It seems that the deductive moves from $A \land B$ to A, or from $A \rightarrow B$ and A to B, are uninformative. All such moves seem utterly trivial. The problem then is that any deductive inference can be reconstructed by chaining together enough of these seemingly trivial inferences. If each step is trivial and uninformative, then we seem committed to saying that the entire deductive inference is trivial and uninformative. Yet some deductive inferences are not trivial, and can be informative. Something is amiss here. 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.
>
> (Dummett 1978a, 297)

(p.197) According to Jago (2013b), the problem has the structure of a sorites series. Suppose you've just marked 100 student essays and, as it happens, each got a different percentage mark from all the others. (So, each positive integer up to 100 is the grade of exactly one of the essays.) The top-marked ones were great. The lowest marked ones were pretty awful. But it's hard to say precisely which ones were good, which ones not good. Is it that all and only those with a mark over 40%, or 55%, or 68%, were the good ones? If so, what about the essay which scored 40% (or 55%, or 68%)? Was it so much worse that the essay which scored just 1% more? Surely not!

It seems absurd to pronounce for sure that only those essays scoring over (say) 55% were any good. Since there's no appreciable difference in quality between each percentage point, it seems that, if we judge any essay to be good, we should also judge the one scoring just 1% less to be good, too. But, as the 100% essay is clearly good, we are then at risk of judging, incorrectly, that all are good. The puzzle is to make sense of truth and inference in a vague language, so that not every essay is counted as being good.

Similarly, the task in the case of deduction is to make sense of a notion of content such that some, but not all, valid deductions are informative. And just as in the case of the essays, we have to do this without drawing an artificially sharp line between those deductions that are informative and those that are not. On this way of thinking about things, the normative notion of logical content is a vague notion, because chains of seemingly uninformative inferences can give rise to informative deductions.

Saying that the content of logical inferences may be indeterminate is not to provide a solution to these issues, however. It is merely to indicate that the problem has a certain form, one which we meet in other cases of vague predicates. Nor is this to say that we can pass the buck, by placing the problem of logical information at the feet of those working on theories vagueness in general. A philosophical theory of vagueness, as commonly understood, is a theory of how vagueness arises (is it metaphysical? semantic? epistemic?) together with an account of how vague predicates work. If we agree that **(p. 198)** the language of logical information, content, and inference can be vague, then a full solution will certainly need to appeal to a general philosophical theory of vagueness. But a full solution to our present problem requires more than this.

A solution to our problem should consist in a model of logical information which explains why we find trivial inferences utterly uninformative, yet capable of being chained together into informative deductions. Let's consider further the analogy with more common cases of vagueness. In a deduction-sorites, each inference rule may be associated with a tolerance principle, saying that if such-and-such deduction is uninformative, then so is the one extended in such-and-such way.

If we have a trivial derivation of A from premises Γ, for example, then the tolerance principle associated with Disjunction Introduction says we also have a trivial derivation of $A \vee B$ from Γ. If we write the relationship of trivial derivation as '$\vdash_{\text{triv}}$', then this tolerance principle can be written:

$$\frac{\Gamma \vdash_{\text{triv}} A}{\Gamma \vdash_{\text{triv}} A \vee B}$$

There are similar tolerance principles for each connective, covering both appearances of the connective on the right-hand side (as the conclusion) and on the left-hand side (in the premises).

Together, these principles give us a proof system for $\vdash_{\text{triv}}$, which coincides with the underlying derivability relation. In other words, $\Gamma \vdash_{\text{triv}} A$ iff $\Gamma \vdash A$: all derivations are trivial! Since that's clearly wrong, at least some of these tolerance principles (expressed as proof rules) are incorrect. Logic dictates that a solution must reject the tolerance principle for at least one connective in each functionally complete set (such as $\{\neg, \wedge\}$, $\{\neg, \vee\}$, and $\{\rightarrow, \bot\}$). If we did not, we could infer that all derivations are uninformative. So one option may be to reject the tolerance principles for some (e.g., ¬ and →) but not all connectives.

We claimed in Jago (2014b), however, that *all* of these tolerance principles should be rejected. The argument is that the inference **(p.199)** rules for '∧' and '∨' stand to the meaning of those concepts just as the inference rules for '→' stand to its meaning. (That is not to say that those rules constitute those meanings, but merely that there is a clear relationship between meaning and inference rules.) So, if the meaning of '→' does not guarantee that uses of *modus ponens* are uninformative, then neither can the meanings of '∧' and '∨' guarantee that inferences involving '∧' and '∨' are uninformative. But what could guarantee that a given kind of inference is always uninformative, if not the meanings of the logical terms involved?

That, in short, is the case for thinking that each of these tolerance principles should be rejected. As a consequence, any inference (other than from a sentence to the very same sentence) might be informative. But it does not follow from this that all inferences are informative. As in the case of other tolerance principles, it is likely that most instances are true, even though the universal generalization is false.

To solve paradoxes involving vagueness, it is not enough merely to reject tolerance principles. One has to explain why they seem so tempting in the first place. (And in deduction-sorites cases, it seems, the tolerance principles are especially beguiling.) In general, we might hold that tolerance principles are false but with a very high degree of truth; or that they are false but any counter-instances are unknowable, and hence unassertable; or that counter-instances shift from precisification to precisification, and so cannot determinately be recognized. Whichever explanation we use, a structure is required which preserves what Fine (1975b) calls the *penumbral connections*. In our case, if an inference is determinately informative, then any inference which includes it is also determinately informative. (And hence, if an inference is determinately uninformative, then any inference which it includes is also determinately uninformative.)

Jago's (2013b, 2014a, 2014b) models use proof rules as links between worlds, rather than as closure principles on worlds. To simplify somewhat: a proof rule directly connects a world where all the premises but not the conclusion is true to a world that's exactly the same, except that the conclusion is true too, according to that world. This connection is directed, from the 'premise' world to the **(p. 200)** 'conclusion' world. Rules with two premises connect two premise worlds to a conclusion world. (Since the worlds in question cannot be logically closed, they are all impossible worlds. But they need not be inconsistent: they could be consistent but incomplete.)

If our proof rules are taken from the sequent calculus, then there's a very direct relationship between proof rules and world-connections. To each world *w* we can associate two sets of sentences, $|w|^+$ and $|w|^-$: those that are true, according to *w* and those that are false, according to *w*, respectively. Then each sequent rule of the form

$$\frac{\Gamma_1 \vdash \Delta_1}{\Gamma_2 \vdash \Delta_2}$$

generates a connection from w_2 to w_1 when $|w_1|^+ = \Gamma_1 \cup \Gamma_2$, $|w_1|^- = \Delta_1 \cup \Delta_2$, $|w_2|^+ = \Gamma_2$, and $|w_2|^- = \Delta_2$. (Note how the connection goes from lower to upper sequent, for this is how, in practice, sequent proofs are constructed.) Rules with two upper sequents generate two of these connections (from the lower to each of the upper sequents).

Chaining these connections together gives a connected graph on worlds. That total graph is a *tree* (that is, a connected, acyclic graph, so that any two worlds are connected by exactly one path). Some of its *subtrees* (those parts of the whole graph that are themselves trees) correspond to proofs. That happens for a subtree *T* when three conditions are met:

(9.1) For each *leaf-world w* (the world found at the end of some branch) of *T*, $|w|^+$ and $|w|^-$ overlap (so that some *A* is a member of both);

(9.2) Every non-leaf node of *T* has at most two edges leading away from it; and

(9.3) There are edges $\langle\langle w_1, w_3\rangle\rangle$ and $\langle\langle w_2, w_3\rangle\rangle$ in *T* only if the proof system contains a rule-instance:

$$\frac{|w_1|^+ \vdash |w_1|^- \qquad |w_2|^+ \vdash |w_2|^-}{|w_3|^+ \vdash |w_3|^-}$$

(p.201) These subtrees are *world proofs*. In effect, they uncover any hidden contradictions in an inconsistent and incomplete world, by connecting it to blatantly inconsistent worlds, according to which some *A* is both true and false. Such blatantly inconsistent worlds aren't epistemically possible for any agent, and so can’t play a role in our notion of epistemic content. But other inconsistent worlds may do so, if their inconsistencies are buried deeply enough. (This raises an important worry: don’t some people, rightly or wrongly, believe contradictions? We’ll defer our discussion until §10.6.)

Our proposed epistemic possibility condition goes like this:

> (EP) World *w* is epistemically possible just in case *w* isn’t the root of any small world-proof.

This is an *absolute* notion of epistemic (im)possibility. If *w* is the root of any small world-proof, then it is *deeply* epistemically impossible, and not just epistemically impossible for some agent or other. Deep epistemic impossibilities are not eligible for playing a role in epistemic notions of content, and so cannot figure in our account of the content of a deduction. All the worlds not ruled out by this criterion are *deeply epistemically possible*, and together constitute *epistemic space*. (We’ll discuss epistemic space in more detail in §10.3 and §10.4.)

‘Small’ is a vague concept. That is ultimately where vagueness enters our account of information content. If it is indeterminate whether *w* is the root of some small world-proof, then it is indeterminate whether *w* is an epistemically possible world, and hence indeterminate whether *w* may play a role in any content. If *A* is true according to *w*, then it will be indeterminate whether *w* is a member of *A*’s content.

Let’s see how this is supposed to help with the problems of logical information and informative inference. In Jago 2013b, the content of *A* is analysed as a pair of sets of worlds: those according to which *A* is true, $|A|^+$, and those according to which *A* is false, $|A|^-$. Call these sets the *positive* and *negative* contents of *A* (so that a content as a whole is a pair of a positive and a negative content). For sets Γ, we have:

(p.202)

$$|\Gamma|^+ = \bigcap_{A\in\Gamma} |A|^+ \qquad\qquad |\Gamma|^- = \bigcup_{A\in\Gamma} |A|^-$$

And finally, the content of a deduction from premises Γ to conclusion *A* is defined as the set of all epistemically possible worlds according to which Γ is true but *A* is false: $|\Gamma|^+ \cap |A|^-$.

As a consequence, the clear cases of trivial inference all come out as contentless and hence uninformative. Take *modus ponens*, from $A \rightarrow B$ and *A* to *B*. Its content is defined as $|A \rightarrow B|^+ \cap |A|^+ \cap |B|^-$. Suppose this set contains a world *w*. Then by definition, there are sets of sentences Γ and Δ such that $|w|^+ = \Gamma \cup \{A \rightarrow B, A\}$ and $|w|^- = \Delta \cup \{B\}$. The sequent rule for → (on the left) has this instance:

$$\frac{\Gamma, A \vdash A, B, \Delta \qquad \Gamma, A, B \vdash B, \Delta}{\Gamma, A \rightarrow B, A \vdash B, \Delta}$$

and so *w* is the root of a world-proof:

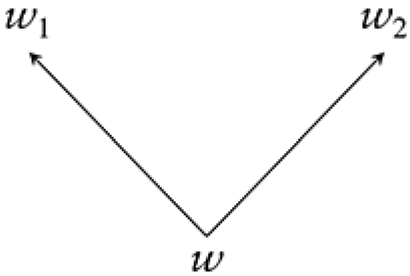

where $|w_1|^+ = \Gamma \cup \{A\}$, $|w_1|^- = \Delta \cup \{A, B\}$, $|w_2|^+ = \Gamma \cup \{A, B\}$, and $|w_2|^- = \Delta \cup \{B\}$. This is a world-proof (and not merely a world-graph) because it is a tree, with leaf-worlds w_1 and w_2, in which $|w_1|^+$ overlaps $|w_1|^-$ (they share *A*) and $|w_2|^+$ overlaps $|w_2|^-$ (they share *B*).

This world-proof is small (by any reasonable standard of proof size). So by definition, *w* is not an epistemically possible world, and hence is not a member of any sentence's content. It follows that *w* it cannot be a member of the content of any inference, contrary to **(p.203)** our original assumption. So $|A \rightarrow B|^+ \cap |A|^+ \cap |B|^-$ is empty, and *modus ponens* is correctly deemed an uninformative inference by our approach. Similar reasoning applies to other clear cases of trivial inferences.

Yet not all valid deductions are deemed empty on this approach. For large *n*, the deduction

$$p_1, p_1 \rightarrow p_2, p_2 \rightarrow p_3, \ldots, p_n \rightarrow p_{n+1} \vdash p_{n+1}$$

is contentful. Its content consists of epistemically possible worlds according to which p_1 and each $p_i \rightarrow p_{i+1}$ ($i < n$) are true but p_n is false. There are infinitely many worlds *w* for which the shortest world-proof with *w* at its root corresponds to $n - 1$ applications of the rule for '→' on the left. (To see this, just consider a world according to which nothing else is true or false, and then consider all the consistent ways of extending *w*.) Since by assumption *n* is large, all such worlds count as epistemically possible, and so the inference is contentful.

We can then formalize the notion of *trivial inference* by taking an inference to be non-trivial just in case it is contentful, in the above sense. Then a valid inference from Γ to *A* will be non-trivial just in case there is some epistemic possibility according to which the premises are true but the conclusion is false. So an inference is trivial just in case every epistemic possibility which represents the premises as being true does not also represent the conclusion as being false.

We will write ‘triv(Γ, A)’ for ‘the inference from Γ to A is trivial’. To formalize the idea, we will need to be precise about which worlds count as epistemic possibilities. The simplest way to do this is to fix, artificially, precisely which world proofs are to count as the small ones. Following Jago (2014b), we pick an integer n, and say that all world proofs of size $m \leq n$ are small. Then, relative to our chosen n, we can say precisely which worlds are epistemically possible, and hence which inferences are trivial. We’ll then write ‘$\text{triv}_n(\Gamma, A)$’ for ‘relative to our chosen n, the inference from Γ to A is trivial’.

Now for the formal details. We use the standard propositional language $\mathscr{L}$ from before, with each connective $\neg$, $\wedge$, $\vee$, $\rightarrow$, and $\leftrightarrow$ as a primitive (undefined) symbol.

(p.204) Definition 9.1 **(Models)** *A model is a tuple* $M = \langle\langle W, N, \rho\rangle\rangle$*, where* W *is a set of worlds,* $N \subseteq W$ *is the subset of normal worlds, and* ρ *is a valuation relation (as in §5.4), relating (i) each atomic sentence to exactly one truth-value at worlds in* N *and (ii) each sentence to zero, one, or two truth-values at worlds in* $W - N$*. A pointed model is a pair,* $\langle\langle M, w\rangle\rangle$ *where* w *is a world in* M*. We abbreviate* $\langle\langle M, w\rangle\rangle$ *to* M^w.

We then extend ρ to all sentences at worlds in N via the standard recursive clauses, as in §5.4. Then for worlds $w \in N$, $\rho_w A1$ iff not $\rho_w A0$, whereas for $w \in W - N$, ρ_w behaves arbitrarily.

Definition 9.2 **(Rank)** *Given a model* $M = \langle\langle W, N, \rho\rangle\rangle$ *and a world* $w \in W$*, we define* w*’s rank,* $\#w$*, as the size (number of nodes) in the smallest world-proof rooted at* w*, if there is one, and* ω *otherwise. The rank of model* M *is* $\min\{\#w \mid w \in W\}$.

Intuitively, a model counts as an epistemic space when its rank is not small. If we select n as our artificial precisification of ‘small world proof’, then only models of rank $r > n$ count as epistemic spaces.

Definition 9.3 **(Trivial consequence)** *For any* $n \in \mathbb{Z}^+ \cup \{\omega\}$*,* A *is an* n*-trivial consequence of* Γ, $\text{triv}_n(\Gamma, A)$*, if and only if, for all pointed models* M^w *of rank* $r > n$*:* $\rho_w B1$ *for each* $B \in \Gamma$ *only if not-*$\rho_w A = 0$.

As a definition of (a kind of) consequence, this definition is rather unusual. This is because trivial consequence is not purely about truth-preservation across all epistemic scenarios. In fact, *no* inference (other than *identity*, $A \vdash A$) is preserved across all epistemic scenarios. Rather, a consequence counts as trivial in the current sense when the truth of the premises guarantees *avoidance of falsity* for the conclusion across all epistemic scenarios.

Because of this, $\text{triv}_n(\Gamma, A)$ behaves as a consequence relation in some ways, but not in others, as the following results highlight. (For proofs, see Jago 2014b.)

(p.205) Theorem 9.1 triv_n *has the following properties, for all* $n \geq 1 \in \mathbb{Z}^+ \cup \{\omega\}$:

(a) $\text{triv}_n \subseteq \text{triv}_{n+1}$: *if* $\text{triv}_n(\Gamma, A)$, *then* $\text{triv}_{n+1}(\Gamma, A)$.
(b) triv_n *is monotonic: if* $\text{triv}_n(\Gamma, A)$ *and* $\Gamma \subseteq \Delta$ *then* $\text{triv}_n(\Delta, A)$.
(c) $\text{triv}_n(\Gamma, A)$ *only if* Γ *classically entails* A.
(d) triv_n *is reflexive*.
(e) $\text{triv}_0(\Gamma, A)$ *if and only if* $A \in \Gamma$.
(f) *For* $n \geq 1$, triv_n *is non-transitive and does not satisfy cut: it is not the case that if* $\text{triv}_n(\Gamma, A)$ *and* $\text{triv}_n(\Gamma \cup \{A\}, B)$ *then* $\text{triv}_n(\Gamma, B)$.

So long as n is not too small, the trivial consequences (so defined) include all the inferences we usually call trivial. Table 9.1 gives some examples, showing the minimal value of n for which triv_n holds. It is not the case that $\text{triv}_3(\{A \vee B, \neg A\}, B)$ holds, for example. To see why, consider a model M containing a single world $w \in W - N$ such that $\rho_w(A \vee B)1$, $\rho_w(\neg A)1$, and $\rho_w B0$ (and that's it for ρ). Then M is of rank 4 and represents the premises as being true, but the conclusion as being false. That's all we need for the inference not to be 3-trivial.

Table 9.1: Some trivial consequences

$\text{triv}_2(\{A \wedge B\}, A)$	$\text{triv}_3(\{A, B\}, A \wedge B)$
$\text{triv}_2(\{A\}, A \vee B)$	$\text{triv}_4(\{A \vee B, \neg A\}, B)$
$\text{triv}_3(\{A \to B, A\}, B)$	$\text{triv}_5(\{A \to B, \neg B\}, \neg A)$
$\text{triv}_7(\{\neg(A \wedge B)\}, \neg A \vee \neg B)$	$\text{triv}_8(\{\neg(A \vee B)\}, \neg A \wedge \neg B)$

The notion of trivial inference is interesting in its own right, but it also plays an important role in our account of fine-grained epistemic and doxastic states. We'll return to the idea in §10.5.

(p.206) 9.6 What Is Said

We've been investigating the notion of informative logical reasoning. This is one useful notion of content: the information contained in a valid deduction. But it clearly isn't the only useful notion of information content. Here's another: the information conveyed in a speaker's *saying that such-and-such*. This is the concept of *what is said* in that utterance. There are good reasons for thinking that this notion of content is distinct from the notion we've just been discussing.

To get a handle on what we have in mind by 'what is said', consider how:

> If someone wants to say the same today as he expressed yesterday using the word 'today', he must replace this word with 'yesterday'. … The case is the same with words like 'here' and 'there'.
>
> (Frege 1956, 296)

Here, Frege is making the point that different words can be used to say the same thing. We can 'say the same thing', but at different times, by using 'today', 'tomorrow', 'yesterday', and so on, depending on what day we speak. And by the same token, we can use the very same form of words, in different contexts, to say different things. When you utter the word 'I', you say something about yourself; but when Franz or Mark uses 'I', they say something about Franz or Mark.

This all goes to show that the concept we have in mind, *what is said*, should not be conflated with utterances (or sentence tokens), nor with sentence types. What a speaker *says* is distinct from the words she utters. We use the concept of *what is said* in the sense of *what is communicated* in making a particular utterance, as opposed to the particular way in which that content is communicated.

Saying the same thing by uttering different sentences is not a phenomenon specific to indexicals (such as 'here' and 'now') or to other context-sensitive words. Even with all contextual factors accounted for, we can still say the same thing in different ways. Just consider:

(p.207)

(9.4a) It's sunny and hot today.

(9.4b) It's hot and sunny today.

Suppose Anna and Bec utter these in a conversation on the same day, in the same place (and with the same conversational standards for 'sunny' and 'warm' in play). Intuitively, it seems they are saying the same thing as one another. It would be bizarre to interpret Bec as disagreeing with Anna. Try this. Imagine Bec had instead replied with, '*no*, actually it's warm and sunny today'. That response would be so bizarre, we would have to interpret her meaning using other conversational clues: perhaps as wanting to emphasize the day's warmth. Either way, Bec *says the same thing* as Anna.

This is a *purely logical* case of same-saying. Anna and Bec say the same thing because they use 'and' to connect two predicates, and (we suggest) the order in which terms flank 'and' doesn't affect what is said. We'll offer more examples like this in a moment. The immediate point here is that *same-saying* is not a phenomenon generated only by context-sensitive terms (as the initial quote from Frege may have suggested).

Note that the correct explanation is *not* that speakers of (9.4a) and (9.4b) say the same thing because those sentences are logically equivalent. This explanation would have it that all logically equivalent utterances say the same thing as one another. But this is not the case, as the following example shows:

(9.5) The Liar is both true and false.

(9.6) Claims about large cardinal numbers are neither true nor false.

These utterances do not say the same thing, even though they are (classically) equivalent. It may be that (9.5)'s speaker is a dialethist, such as Priest (1979, 1987), who diverges from classical logic in rejecting the Explosion Principle (that contradictions entail arbitrary conclusions), whereas (9.6)'s speaker is a mathematical intuitionist, such as Dummett (1978b, 1993b), who rejects Excluded Middle. **(p.208)** Each of these philosophical positions is completely different from the other. It is absurd to think that, in stating their different philosophical positions, they say the same thing as one another. So it is not the case that, in uttering any two classically equivalent sentences, the speakers thereby say the same thing as one another. *What is said* is a hyperintensional notion of content.

Yet, as we saw with the example pair (9.4a) and (9.4b), the 'anything goes' approach to content we mentioned in §8.4 doesn't give an appropriate analysis of same-saying. Some logical relations (including the one relating $A \wedge B$ to $B \wedge A$) preserve same-saying. Here are two further pairs in which, we think, an utterance of (a) says the same as an utterance of (b):

(9.7a) Anna or Bec will pass, and Cath will pass too.

(9.7b) Either Anna and Cath will pass, or else Bec and Cath will.

(9.8a) Either Cath doesn't like Dave or she doesn't like Ed.

(9.8b) Cath doesn't like both Dave and Ed.

These strike us as clear examples of same-saying. These same-saying pairs suggest that commutativity for 'and', distributivity of 'or' over 'and', and the De Morgan equivalences, are all operations which preserve same-saying. Replacing 'and' with 'or' in (9.4) also seems to preserve same-saying, so we may add to the list commutativity for 'or'. And associativity for both 'and' and 'or' seems obviously to preserve same-saying.

The following pair is perhaps more contentious:

(9.9a) Valeria is happy.

(9.9b) It is not the case that Valeria isn't happy.

We think an utterance of either says the same as the other. Intuitionists will disagree (at least for certain cases involving double-negations). And to be sure, there may a difference in the meaning conveyed in **(p.209)** English between 'Valeria is happy' and 'Valeria is not unhappy'. But the difference is a shade of meaning, not a difference in literal content; and anyway, we don't find this difference present when using 'it is not the case that', as in (9.9b). So, we say, introducing or eliminating double negations preserves what is said.

We may think of these transformations algebraically, in terms of operations $\wedge$, $\vee$, and $-$ on contents. Then we can write our principles as identities between same-saying contents: $c_1 \wedge c_2 = c_2 \wedge c_1$, $-(c_1 \wedge c_2) = -c_1 \vee -c_2$, $--c = c$, and so on. The examples suggest that we have a *De Morgan algebra*: a *bounded distributive lattice*, with top element W (the set of all worlds) and bottom element $\varnothing$, where $-$ is an involution which obeys the De Morgan laws (see, e.g., Balbes and Dwinger 1975). But we don't have a full Boolean algebra, since we don't in general have $c \vee -c = W$ or $-(c \wedge -c) = \varnothing$.

We can then use what we know of such structures to generate further predictions about same-saying. One is that both $\wedge$ and $\vee$ are idempotent: $c \wedge c = c \vee c = c$. And this seems intuitively right to us:

(9.10a) Valeria is happy.

(9.10b) Valeria is happy and Valeria is happy.

(9.10c) Valeria is happy or Valeria is happy.

all seem to say the same thing as each other (although the use of (b) and (c) would call for a rather strange context, and so may pragmatically convey more than an utterance of (a) alone).

Another consequence of this approach is that $c_1 \wedge (c_1 \vee c_2) = c_1 \vee (c_1 \wedge c_2) = c_1$, and so it predicts that each of

(9.11a) Bertie is snuffling, and either he's snuffling or Lenny is barking.

(9.11b) Either Bertie is snuffling, or he's snuffling and Lenny is barking.

(9.11c) Bertie is snuffling.

(p.210) says the same as the others. But this strikes us as a poor prediction: (9.11a) and (9.11b) both say something about both Bertie and Lenny, whereas (9.11c) says something about Bertie only. So, it seems, (9.11c) can't say the same thing as the others. If that's right, then the algebra in question is weaker than a distributive lattice.

Now let's bring the discussion back to impossible worlds. If we are to understand these same-saying contents as sets of worlds, then we must restrict the worlds in question to those that obey these operations. If $A \wedge B$ is true according to one of the worlds in question, then $B \wedge A$ must be too (and so on for all the other principles just discussed). To put things another way, the logic generated by the algebra in question gives us closure conditions on worlds. Worlds which do not meet those closure conditions are not eligible for inclusion in same-saying contents. Those worlds still exist; it's just that we ignore them when we theorize about same-saying. Whatever we deem the right algebra for same-saying contents, we can model it using fine-grained worlds.

There is an objection to this approach, however. It seems that the notion of a world, in and of itself, is doing little theoretical work in this approach. All the work is done by selecting the right algebra. For given that algebra, it doesn't really matter what kind of entity the c_1s and c_2s are. Formally, they could be sets of root vegetables, and the approach would work just as well as if they were sets of worlds.

The objection then continues: surely a better approach is to select some kind of conceptual tool which *generates* the appropriate algebra, rather than being generated by it? Jago 2018b presents an alternative approach, on which same-saying contents are *sets of truthmakers*. Then two utterances say the same as each other just in case whatever would make one true would also make the other true, too. On that approach, the logic of same-saying equates to *strict truthmaker equivalence* (Fine and Jago forthcoming). This generates all the same-saying equivalences we viewed positively above, but does not imply that either $A \wedge (A \vee B)$ or $A \vee (A \wedge B)$ says the same as A.

We can always mimic this approach using impossible worlds, by focusing on only those worlds which are closed under strict truthmaker **(p.211)** equivalence. The objection, however, is that understanding contents in terms of truthmakers generates the right results automatically, without having to impose further restrictions (in the form of closure conditions) on worlds. In that sense, it might be said to afford us a better understanding of same-saying.

We'll consider two responses on behalf of the impossible worlds approach. The first claims that the truthmaker approach doesn't (clearly) give the best results after all. In strict truthmaker logic, $(A \vee B) \wedge (A \vee C)$ does not entail $A \vee (B \wedge C)$, and so distribution of ∨ over ∧ is not an equivalence. Yet utterances of the following pair seem to say the same as one other:

(9.12a) Either Anna will pass, or both Bec and Cath will pass

(9.12b) Anna or Bec will pass; and also, Anna or Cath will pass.

We'll grant that it's hard to come to a clear view on this. But if we agree that they do say the same thing, it follows that strict truthmaker equivalence isn't the logic of same-saying.

The second is a metaphysical response. Contents exist, and hence we may existentially quantify over the constituents of those contents. If those constituents include possible states of affairs, then we seem to be saying that merely possible states of affairs genuinely exist. (This is the stance of the genuine realist about merely possible worlds, applied to states of affairs.) This cannot be right. If Franz is in Amsterdam, he's not in Italy. Both states of affairs are possible, but at most one exists. If both existed, then Franz would be both fully in Amsterdam and fully elsewhere. That's impossible. We shouldn't accept a semantic theory which requires reality to be like that. There are a number of responses available. One could go in for ersatz states of affairs, or non-obtaining states of affairs, or non-existent states of affairs. But all of these face issues of their own. (See Jago 2018b for an in-depth discussion of the options.)

A positive reason for analysing same-saying contents in terms of possible and impossible worlds, rather than in terms of possible states of affairs, is that it will then integrate well with other notions **(p.212)** of content (such as the information contents of §9.5 and doxastic contents of §10.5). All such contents are defined on the domain of worlds. This allows us to speak of one content including, overlapping, or being disjoint from another, even when they are different kinds of content (informational and same-saying, say). This is important when we want to investigate, say, how the content of what a trusted speaker says affects what their hearer thereby comes to believe.

Chapter Summary

We conceptualize information in terms of ruling out scenarios (§9.2). We discussed informative identity statements, which give rise to Frege's puzzle (§9.3), and the problem understanding how a valid logical inference can be informative (§9.4). We gave an analysis of informative logical inferences in §9.5, on which the content of a valid deduction is often indeterminate. A consequence is that it is indeterminate exactly which logical inferences are informative. We then analysed a rather different notion of content, concerning *what is said* by a speaker in making an utterance (§9.6).