Erkenn (2012) 77:167–185
DOI 10.1007/s10670-011-9356-1

ORIGINAL ARTICLE

Informational Semantics as a Third Alternative?

Patrick Allo · Edwin Mares

Received: 23 July 2010 / Accepted: 10 December 2011 / Published online: 28 December 2011

Abstract Informational semantics were first developed as an interpretation of the model-theory of substructural (and especially relevant) logics. In this paper we argue that such a semantics is of independent value and that it should be considered as a genuine alternative explication of the notion of logical consequence alongside the traditional model-theoretical and the proof-theoretical accounts. Our starting point is the *content-nonexpansion* platitude which stipulates that an argument is valid iff the content of the conclusion does not exceed the combined content of the premises. We show that this basic platitude can be used to characterise the extension of classical as well as non-classical consequence relations. The distinctive trait of an informational semantics is that truth-conditions are replaced by information-conditions. The latter leads to an inversion of the usual order of explanation: Considerations about logical discrimination (how finely propositions are individuated) are conceptually prior to considerations about deductive strength. Because this allows us to bypass considerations about truth, an informational semantics provides an attractive and metaphysically unencumbered account of logical consequence, non-classical logics, logical rivalry and pluralism about logical consequence.

P. Allo (✉)
Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel, Brussel, Belgium
e-mail: patrick.allo@vub.ac.be

P. Allo
IEG, University of Oxford, Oxford, UK

E. Mares
School of History, Philosophy, Political Science and International Relations,
Victoria University of Wellington, Wellington, New Zealand
e-mail: edwin.mares@vuw.ac.nz

1 Introduction

The *prima facie* case for considering "informational semantics" as an alternative explication of the notion of logical consequence alongside model-theoretical and proof-theoretical ones is easily summarised. Where model-theory and truth-conditional semantics are standardly associated with a defence of classical logic (CL), and proof-theory and inferential semantics with a defence of intuitionist logic (IL) as well as other constructive logics like linear logic, information-conditional semantics seems to be wedded to relevant and other substructural logics (RL).[1] As such, if the CL, IL, RL trio is a representative chunk of a broader range of logical options (we deliberately ignore the more general approach of linear logics, especially the non-distributive ones),[2] information-conditional semantics surely has its place. Yet, it is even easier to dismiss the suggestion that information-conditional semantics provides an apparently missing third conception of logical consequence. After all, isn't it just a variant of the usual interpretation of the Routley-Meyer relational semantics rather than a genuine alternative to a model-theoretic account? Or worse, isn't it a mere metaphor? In the present paper, we want to consider a more subtle answer to the question of whether information-conditional semantics is a real alternative for the two more traditional contenders.

Let us, to begin with, recall where so-called informational semantics come from. Two aspects are relevant here; one positive, the other rather defensive. On the positive side, we have the family resemblance between the use of partial information-states in relevant logic—especially its first-degree fragment—and Barwise and Perry's situation semantics (Dunn 1976). The defensive side, by contrast, is to be found in the use of situation semantics, and in particular in the use of the notion of information-flow, as part of the defence of the Routley-Meyer relational semantics—the infamous ternary relation—against the charge that it didn't constitute a genuine semantics. To that end, Restall (1995) used a version of Barwise's channel-theory, and Mares (1997) a theory of information due to Israel and Perry (1990). This resulted in both cases in an interpretation of a pre-existing formalism. More exactly, it extends the analogy between partial information-states and situations beyond the first-degree fragment by including connections between situations, and provides an intuitive picture for the less palpable model or truth-theoretic account of the meaning of relevant implication.[3]

[1] We use the terms "informational semantics" and "information-conditional" semantics interchangeably. The former is appropriate because of the analogy with informational content, whereas the latter emphasises the contrast with truth-conditional and proof-conditional semantics.

[2] Although this restriction is primarily for reasons of focus, a more principled exclusion of linear logic can be mounted by, for instance, denying that linear logic (without distribution) should be seen as a formal model of logical consequence in natural language. This seems consistent with Girard's original motivations for developing linear logic; which he didn't conceive as a rival to classical logic. See Paoli (2007) for a dissenting opinion.

[3] Other approaches could equally well be described as *informational*, e.g. the informational interpretation of substructural logic defended by Wansing (1993a, b), the informational reading of sequents mentioned in Paoli (2002, 33), and the Boolean informational conception described by John Corcoran (see e.g. Saguillo 2009).

If we want to argue that information-conditional semantics constitutes a genuine alternative, it needs to be argued for on independent grounds. This means, in the first place, that we need positive, non-defensive reasons for introducing it. But that is not sufficient. Even if we can make the interpretation plausible enough, it will still be that: a mere interpretation of a pre-existing formal account. To cope with that aspect, we either need to explain why a separate formal account isn't required, or we need to provide such an additional formalism. In this paper, we pursue the former option.

Our discussion undoubtedly leaves many questions unanswered. We mainly try to give the reader an idea of why information-conditional semantics is a genuine and indeed an attractive alternative. To that end, we sketch two complementary pictures of the informational approach to logical consequence: a traditional model-theoretic one, and a more abstract one based on the inverse relation between logical discrimination and deductive strength. The scope of our argument remains nevertheless limited. First, because we only provide a formalism that is sufficiently flexible to include logics that have a relational semantics. This excludes several substructural logics that, like linear logic, do not validate the usual distribution principles. Second, because we consider the whole issue of informational semantics from the perspective of relevant logics. There is nothing essential about these restrictions. We're entirely in favour of a more general account that would subsume even more logics, but we also believe that the trio of classical, intuitionist and relevant logics is sufficiently representative for our purposes. Also, while an informational semantics for relevant logics is indeed our main focus, we also explain how this extends to intuitionist and classical logic.

2 What Is an 'Informational' Semantics?

If we want to properly situate the informational conception of logical consequence, it is best to start with the observation that the orthodoxy has it that a model-theoretical semantics gives a truth-conditional analysis of consequence, and a proof-theoretical semantics an inferential (or proof-conditional) analysis. This is best understood as follows. Reading it from left to right, we might say that the truth-conditional and inferential conceptions of logical consequence are formalised by, respectively, model-theoretic (truth-in-a-model) semantics and by proof-theoretic semantics. When read from right to left, we could then say that model and proof-theoretical semantics can be understood or interpreted in terms of, respectively, the truth-conditional or inferential conceptions of consequence (and, relatedly, the inferential and truth-conditional accounts of the meaning of logical connectives). This means we have a two-way interaction between the formal and the intuitive conception of logical consequence; something that roughly corresponds to the interaction between pure and applied semantics where the former formalises the latter or the latter interprets the former. This is also the terminology in which the debate on the Routley-Meyer semantics was initially phrased (Copeland 1979).

Quite like the two traditional conceptions, the informational conception is based on a platitude about what it means to "follow from," namely the suggestion that the content of a conclusion should not exceed the joint content of the premises.

CN A conclusion A follows from premises Γ iff the content of A does not exceed the combined content of all the premises in Γ.

However plausible, this description is problematic for several reasons. First, one could simply claim that **CN** is precisely what needs to be explicated. **CN** is either too generic to be useful, or merely a reformulation of the intuitive notion of consequence. Relatedly, the notion of content is often thought to be parasitic on the notion of meaning, but one thing a conception of logical consequence should explain is precisely the meaning of the logical connectives. Finally, if we understand **CN** as saying that the relative semantic content of A given Γ should be null, and formalise this along the lines set out in Carnap and Bar-Hillel (1952) and Kemeny (1953), this leads us straight back to a model-theoretic characterisation. In other words, **CN** is apparently just truth-conditional semantics in disguise.

These objections are best dealt with in reverse order, starting by admitting that a Carnap-style formalisation of semantic content is not directly helpful if we want to explain the distinctive features of informational semantics. More exactly, if we stipulate that A follows from Γ iff $\mathrm{CONT}(A|\Gamma)$ is null, an examination of the following definition of relative content immediately reveals the underlying model-theoretic framework.

$$\begin{aligned}\mathrm{CONT}(A|\Gamma) &= \mathrm{CONT}(\{A\} \cup \Gamma) \setminus \mathrm{CONT}(\Gamma) \\ &= \{w : w \Vdash \Gamma\} \setminus (\{w : w \Vdash \Gamma\} \cap \{w : w \Vdash A\})\end{aligned} \tag{C}$$

with w ranging over the set of possible worlds, state descriptions, or some other class of complete and consistent *cases*.

Note that this doesn't yet mean that a broadly Carnapian implementation of **CN** is per se also a truth-theoretic explication of consequence. That's not the issue here. Instead, the complaint is that because it is tied to a classical conception of content, there is no way to distinguish an information-conditional from a truth-conditional account of consequence. This is because (C) adopts a worldly perspective; which comes down to the following two features. First, it already integrates the structural features of classical logic. Second, and more importantly, (C) analyses content in terms of an extensional notion of *exclusion*: The content of A are the non-A cases.

On the conceptual level, (C) may be judged inadequate because an extensional analysis of informational content reduces the content-nonexpansion account of logical consequence to a truth-conditional account. This is illustrated by the formal inadequacy of (C), and especially by the fact that it cannot be used as a schema to obtain non-classical accounts of informational content. If we assume that s ranges over the set of possibly incomplete and/or inconsistent situations, we do not have that $\mathrm{CONT}(A|\Gamma) = \{s : s \Vdash \Gamma\} \setminus (\{s : s \Vdash \Gamma\} \cap \{s : s \Vdash A\})$ is null iff A is relevantly entailed by Γ. This is a clear failure to comply with the content-nonexpansion platitude.

To generalise (C) to include the relevant case, we need to use the following definition (Allo 2007, 682):

$$\begin{aligned} \mathrm{CONT}_c(B|A) &= A_c^{\sqsubseteq} \setminus (A_c^{\sqsubseteq} \cap \{s : s \Vdash B\}) \\ &= A_c^{\sqsubseteq} \setminus \{s : s \Vdash A \otimes B\} \end{aligned} \tag{R}$$

with $A_c^{\sqsubseteq}$ defined in terms of the ternary accessibility relation $\sqsubseteq_{\ldots}$ over the set of situations S (i.e. we write $s \sqsubseteq_c s'$ instead of the more common notation $Rcss'$).

$$A_c^{\sqsubseteq} = \{s : \exists t(t \Vdash A \,\&\, t \sqsubseteq_c s)\},$$

and $\otimes$ as the *fusion*-connective with the following satisfaction-clause:

$$s \Vdash A \otimes B \text{ iff } \exists t, u \in S \text{ where } t \sqsubseteq_u s, t \Vdash A, \text{and } u \Vdash B$$

Unlike (C), this new definition can be used as a schema. When s ranges over the set of situations, we obtain a relevant account of content; when it ranges over the set of possible worlds, (R) can be shown to reduce to (C). This formal point reinforces our informal remark that as far as (C) is concerned, there is no difference between explaining consequence in terms of content-nonexpansion (i.e. the requirement that $\mathrm{CONT}(A|\Gamma) = \emptyset$) and explaining it in terms of truth-preservation (see Hanson (1980) for a similar diagnosis). Similarly, as suggested in Beall et al. (2011, 3.2), when we identify the first two indices in the ternary relation (and make suitable modifications to ensure the proper behaviour of negation), (R) can be used to individuate intuitionist information because s then ranges over constructions.[4] In the next section we briefly investigate how this notion of constructive content reduces to proof-conditions. For now, we proceed by focusing on the contrast between truth-conditions and information-conditions.

Conversely, where (C) could still be explained in truth-theoretic terms, namely by reading $w \Vdash A$ as A is true at w, this isn't as easy when we formalise **CN** along the lines of (R). Remember: We don't want information-conditional semantics to be a type of truth-conditional semantics in disguise; and, as we shall stress later on, we don't think that the non-classical properties of logical consequence should automatically carry over to our theory of truth. What we need is a reading of $s \Vdash A$ as an information-condition, but this reading can itself not be explained in terms of (R) itself because (R) is an account of informational *content* and not of information.[5] In other words, we need a different, more primitive, account of information.

This is where situation semantics come in. If we are prepared to adopt an information-conditional theory of meaning, we can side-step the objection that an account of meaning is already presupposed by an account of content. Consequently, all we need to do to make **CN** respectable is to complement it with a suitable notion of information; one that can be used in combination with (R). The point of formulating an information-conditional semantics is not just to replace "truth" by

[4] To see why this holds, just note that for consistent situations $s \sqsubseteq_c t$ reduces to the information-containment relation $s \sqsubseteq t$ that is familiar from the Kripke-style semantics for intuitionist logic. We do not pursue this in further detail, but will come back to the issue of intuitionist content in relation to proof-theoretical semantics and the inferential account of consequence.

[5] See also MacKay (1969, 56) on the need to distinguish between the measure of a thing and the thing itself. This worry obviously extends to qualitative individuations of informational content.

"information". Giving a reading of $s \Vdash A$ as an information-condition presupposes an inversion of the order of explanation: Information is used to explain meaning as well as consequence.

As already conceded in Mares (2010), situation semantics belongs to the model-theoretic tradition, but it differs from the mainstream by identifying meaning with information-conditions, *not* with truth-conditions. This explains why information-conditional semantics might not need a dedicated formalism to count as a genuine alternative account, and why it primarily needs to supply a different reading of expressions of the form $s \Vdash A$ that can be plugged into (R). Even so, the deployment of information-conditional semantics in terms of situated inference (Mares 2004) isn't free of proof-theoretical considerations either, but is the result of a dialogue between proof-theoretical and model-theoretical considerations about how information is *accessed* and *used*.[6] It is not—and this should clearly be stated—the result of mixing truth-conditional with inferential semantics. Informational semantics do not merely *supplement* the truth-conditional account with an additional inferentially inspired relevance-condition,[7] but it *replaces* truth-conditions as well as proof-conditions and proposes a genuine alternative to both.

The dependence of informational semantics on model- and proof-theoretical considerations cannot always successfully be distinguished from how it is related to truth-conditional and inferential conceptions. One reason is that the difference between proof-theoretical semantics and inferential conceptions of consequence is itself fairly thin. If a distinction is to be made, it will most likely depend on how sequents are read. Deductive systems are neutral on that matter, but the inferential and information-conditional conceptions are not. Another reason is that the partiality of information-states and of situations is often explained in terms of a non-classical truth-theory; one that includes *cases* where truth and falsity appear to behave non-standardly. The latter move—which conflates the information that is available in a situation with what is true in that situation—is understandable, but unnecessary. Not all situations need to be continuous spatio-temporal parts of the world in the strictest sense, they could, as suggested in Perry (1986, 101), equally well be "parts of the world that determine the answers to a certain set of issues" or abstract situations. The containment-relation between worlds and situations is purely based on information-conditions. This means that abstract situations can be part of the world if they contain information about a concrete situation (or several such situations) in that world. In fact, resisting the adoption of a non-classical truth-theory is a good way to avoid a dialetheic interpretation of inconsistent situations (Mares 2008).

At the core of our defence of information-conditional semantics is an inversion of the standard order of explanation. This occurs at two levels. First, the standard connection between meaning and information is inverted. Information doesn't depend on a prior account of meaning (as with Dretske (1999)), but is used to

[6] This means that in some cases, the semantics is justified in terms of a prior inferential practice (Mares 2009a, 348).

[7] E.g. "real use"; see also Read (1988, 2003) for a criticism of the view that relevance is an additional necessary condition we need to impose on logical consequence.

naturalise meaning. This identification of meaning and information is familiar from ecological realism as well as from situation semantics (see e.g. Turvey and Carello (1985) for a comparison). We won't deal with this aspect here, but just note the crucial inversion. Second, information is no longer explained in terms of the notion of logical possibility (classically, this means in terms of consistency and inconsistency) as we find it in Carnap and Bar-Hillel (1952) where informational content is analysed as the exclusion of logical possibilities (in fact, consistent state-descriptions), but directly in terms of how information is *accessed* and *used* (the precise connection will be explained later on, but see Barwise (1997) for a similar inversion of the standard order of explanation). This inversion is more directly relevant for the purpose at hand. In particular, it can be used to explain the difference between the use of (R) and the use of (C) in the identification of valid consequence with null relative content. While both are formulated extensionally (using set-theory), the latter depends on a prior notion of logical possibility, whereas the former encodes what it means to access and use information that is available in an environment. As a result, one explains informational content in terms of logic, but the other explains logic in terms of informational content.

3 Access and Use

The idea is that the notion of information or informational content which figures in **CN** is explicitly restricted to objective information that is available in a given environment.[8] That is, for A to follow from Γ the information that A may not exceed the information that is accessible in any environment where the information that Γ is accessible as well. The relation between environments and access, and different consequence relations is closely related to Beall and Restall's argument for logical pluralism (2006). A classical consequence relation is obtained by identifying accessible information with truth. This corresponds to environments as possible worlds, and information-states as sets of such worlds. Relevant and intuitionist consequence relations are obtained by further qualifying the relation between truth and information; either by identifying accessible information with the information contained in a situation for the former, or by identifying it with the information contained in proof-stages and extensions thereof for the latter.

By comparing truth with the classical, relevant, and intuitionist notions of information we get a better understanding of how these three logics are related. We have already seen that on the classical account accessible information reduces to truth in a set of worlds. The intuitive picture on which the relevant notion of information is based, is the view that A is true at a world w iff there is a situation s in

[8] For the interpretation of the first-degree fragment of relevant logic, there is no need to be exceedingly precise about the notion of information that is at play. It could equally well be explained in terms of subjective information (the information stored in a database) as in terms of objective information (the information accessible in an environment). Roughly speaking, the former is closely related to the so-called American plan which uses a four-valued semantics for relevant logic (think, for instance, of Belnap's "How a computer should think" Belnap (1976) reprinted in Anderson et al. (1992, §81)), while the latter is tied to the Routley-Meyer semantics that is distinctive for the Australian plan.

w such that s contains the information that A. This is essentially a logical point. It relates the information available in a situation with the extension of truth at a world, but doesn't explain truth in terms of information (or vice-versa). Truth at the world is necessary for a situation to carry that information, but given that situations can be partial and that the information about all that is the case is often distributed, the truth of A at a world w isn't sufficient for *all* situations in that world to contain the information that A. Furthermore, because abstract as well as concrete situations can be part of a world, information-conditional semantics isn't committed to the view that truth in a situation is also necessary for that situation to carry information about the truth in question either. Apart from the necessary connection with truth at a world (Floridi 2005), we don't need to be very specific about the nature of objective information. It may be thought of as a primitive notion, or it may be analysed in more primitive terms (different from meaning and consequence). For instance, information could be conceived as a reliable indicator of truth.

The more urgent matter is what it means for information to be accessible, and how accessible information may be put to use. Access, or rather the lack of access, is closely related to the distributed nature of information. It involves being in the right environment, but also being the right kind of agent. With a slogan: Information is somewhere, but not everywhere, and it is potentially for someone, but not for everyone. Because information is distributed, using information in inference often involves combining information from different situations. This has far-reaching implications for our semantics. We illustrate this by explaining how the notions of access and use are built into (R). The explanation uses the following concepts:

1. The ternary relation $s\sqsubseteq_c s'$ says that combining information that is available in situation s with information available in situation c will, when both situations are in the same world w, provide information about a situation s' that is also in w (but is possibly distinct from s).
2. The set of situations $A^{\sqsubseteq}_{c}$ are those situations about which we can gain information by combining the information that A with information available in c.
3. $A \otimes B$ (read: A *fusion* B) is the result of combining the information A and B.
4. The set of logical situations *Log* are those situations which tell us about constraints that hold across all situations. (We shall henceforth refer to the information present in all logical situations as the *logical information*.)

The intended reading of $\mathrm{CONT}_c(B|A)$ is the relative complement of the set of situations we learn about by combining the information that A with the information that B in the set of situations we may learn about by combining the information that A with information available in c. On that account, B is a logical consequence of A iff the result of combining A and B doesn't tell us anything about a situation we couldn't already learn about by combining the information that A with the information available in a logical situation (i.e. for $c \in Log$).

The above interpretation allows us to illustrate how a space of possibilities can be considered as an abstraction of how we access and use information, rather than as something that is conceptually prior and therefore determines this practice. From a

relevant perspective, access and use are essentially characterised by the following two features: We access information locally, and we combine information from different situations without always having access to a situation that contains the sum of all the information we used. To accommodate the former, the logical space should contain situations that do not decide every issue. This is achieved by partial situations. To accommodate the latter, the situation(s) we learn about by combining information need not contain all the information that was combined. This is achieved by allowing $s \sqsubseteq_c s'$ while $s \not\sqsubseteq s'$, which means that the information in s need not be contained in the information in s'.[9] In the Routley-Meyer relational semantics this reduces to saying that c does not have to be a logical situation. As a result, from a relevant perspective the assumption that inference isn't always cumulative and the assumption that one doesn't have a global access to logical information are two sides of the same coin. In the presence of a *De Morgan* negation, inconsistent situations fall out naturally from these assumptions.

From an intuitionist perspective, access and use have slightly different features: Access is still local, but the process of combining information is cumulative as well as consistency preserving. To accommodate the cumulative nature of inference, the Routley-Meyer semantics stipulate that we need to identify $s \sqsubseteq_c s'$ with $s \sqsubseteq s'$. By definition the latter only holds when c is a logical situation, and thus a cumulative and consistency preserving inferential practice reduces to having global access to all logical truths. Yet, if we have a *De Morgan* negation it is easy to show that logical situations decide every issue.[10] This is undesirable from an intuitionist perspective, as it reduces the global access to logical truths to obtaining a worldly perspective. To allow for logical yet incomplete situations, we need to avoid certain properties of a *De Morgan* negation. Specifically, we should have situations that contain the information that $\neg\neg p$, but not the information that p. This is what characterises constructions in the sense of Beall and Restall (2006).

We have seen that relevant and intuitionist logics formalise two ways of accessing and using information, and that both lead to different ways of carving up the logical space. The intuitionist approach combines local access with (consistency driven) cumulative use via the absence of completed constructions. The relevant approach combines local access with non-cumulative use via the presence of situations that are not logical. The comparison with the practice encoded by classical logic is then straightforward. When compared to relevant logics, the main difference is that classical logic presupposes that one's access to the body of logical information is global (or, equivalently, that the process of combining information is cumulative). When compared to intuitionist logics, the main difference is that

[9] In virtue of $s \sqsubseteq_c s'$ iff $c \sqsubseteq_s s'$, it follow that the information in c also doesn't have to be contained in the information in s'.

[10] Let 1 be propositional constant such that $s \Vdash 1$ iff $s \in Log$. Note furthermore that in a Routley-Meyer semantics where the clause for negation is given by $s \Vdash \neg A$ iff $s^* \not\Vdash A$, we have that (1) $s \sqsubseteq s^*$ holds iff s is consistent, (2) $s = s^*$ holds iff s is maximal, and (3) the validity of double-negation elimination is ensured by $s^{**} = s$. To show that all logical situations are maximal, let s be a logical situation. This can be expressed as $s \Vdash 1$. Furthermore, since all logical situations are consistent, we have $s \sqsubseteq s^*$. Thus we may conclude that $s^* \Vdash 1$, which is just to say that s^* is a logical and *a fortiori* also consistent. Applying the same reasoning to s^* as we did for s, we may conclude that $s^* \sqsubseteq s^{**}$. Because s^{**} is just s, we thus have $s^* \sqsubseteq s$. Since we've already established that $s \sqsubseteq s^*$, we conclude that $s = s^*$ as required.

classical logic presupposes a global access to maximally consistent bodies of information. Either way, the corresponding logical space should only contain complete and consistent situations.[11]

While we believe that the present illustration is illuminating, we do not wish to overemphasise the import of the model-theoretic implementation of the content-nonexpansion platitude given in (R). The above illustration primarily shows how an account of information that is based on a space of possibilities can be seen as an abstraction of how we access and use information. As far as we are concerned, an informational reading of sequents could have done the job as well. In that case, structural rules would have been motivated by an informal understanding of how we access and manipulate information. We do not pursue this path any further, but just point out that structural rules govern our use of assumptions as well as the way we combine premises. The latter can straightforwardly be seen as core features of how we access and use information.

4 Global Constraints and Granularity

In the previous section we explained how a difference in perspective could be used to explain a difference in the information accessible in an environment and hence a difference in the appropriate logic to analyse a given deductive practice (access and use). Here, we consider the matter in a more abstract fashion by focusing directly on the inverse relationship between *logical discrimination* and *deductive strength*. This relationship is a formal feature that, independently of the semantics, holds for many logical systems (Humberstone 2005), and that can be used to elucidate the specificity of informational semantics.

Loose references to the notion of logical discrimination are familiar from how we compare classical with non-classical logics in terms of what can be "told apart" in these logics. For instance, we say that paraconsistent logics (like for instance relevant logics) allow us to discriminate between different inconsistent theories, whereas classical logic, which reduces all inconsistent theories to the trivial theory, cannot. Likewise, we say that intuitionist logic discriminates between p and $\neg\neg p$ while classical logic doesn't.

Depending on the formalism one uses, differences in logical discrimination surface in different guises. For instance, in a relational semantics were propositions are modelled as sets of points, two propositions can be discriminated iff they correspond to different sets of points. Taking these points to be situations, and propositions to be sets of situations that are upwardly closed under the information-containment relation $\sqsubseteq$ yields a finer way of carving out propositions than when points are possible worlds (i.e. complete and consistent situations) and propositions sets of such worlds. This can be seen from the fact that two upwardly closed sets of

[11] The here described collapse of *De Morgan* and *Intuitionist* negation into *Boolean* negation only works on the assumption of distributivity. As described by Dunn (1993), in the absence of distribution we only obtain *Ortho*-negation.

situations may be distinct, and yet coincide in their respective subsets of complete and consistent situations.

A more general account of the same phenomenon can be given by applying Humberstone's preferred criterion for logical discrimination directly. On that account, two formulae A and B cannot be discriminated iff they are synonymous, that is:

$$C_1(A), \ldots, C_n(A) \Vdash C_{n+1}(A) \text{ iff } C_1(B), \ldots, C_n(B) \Vdash C_{n+1}(B)$$

where each $C_i(B)$ is obtained from $C_i(A)$ by replacing some (but not necessarily all) occurrences of A in $C_i(A)$ by B. Using this criterion, we can show that the inability to discriminate between different contradictions in a given logic suffices for that logic to be explosive. By identity we have $p \wedge \neg p \Vdash p \wedge \neg p$; from the assumption that all contradictions are synonymous, we then obtain $p \wedge \neg p \Vdash q \wedge \neg q$. Finally, by *simplification* we have $p \wedge \neg p \Vdash q$ for arbitrary q. This is yet another way to make explicit why "[c]lassical consequence fudges all distinctions among inconsistencies and among tautologies" (Beall and Restall 2006, 31), but relevant logics don't.

Put in the terms of the previous section, the rough idea is that the accessibility or availability of information is relative to logical discrimination: From a worldly perspective we "access" more because propositions are more coarsely individuated (if two tautologies, or two contradictions express a single propositions, we cannot access one without accessing the other). This is a particularly attractive idea for the logical pluralist. If all there is to a difference in logic is a difference in granularity, pluralism about logical consequence follows from the lack of a uniquely correct way of carving out propositional contents.[12]

But in what sense is this different from truth-conditional semantics? Traditional (truth-conditional) treatments of non-classical logics generalise the truth conditions that are found in the standard semantics for classical logic in a way that can equally well be understood as discriminating propositions more finely. The classical truth tables or models for classical first order logic are usually special cases of the models for a given non-classical logic. Non-classical logicians have used this fact to give a classical gloss to their understanding of the truth conditions for the connectives. For example, Dunn's four-valued semantics for first-degree entailments assigns zero or more truth values to a formula (Dunn 1976). The classical tables are generated by those valuations that assign one and only one value to each atomic formula (and hence to all formulas of the language). The truth-conditional interpretation of this semantics holds that sentences can have zero, one, or two truth-values, but that the evaluation of conjunction, disjunction, and negation are very classical. Pluralists can therefore tell two different stories. They can hold that what is different between the classical logician and the four-valued logician is that the latter is willing to consider more *cases* in evaluating truth values and inferences, but they can also claim that a

[12] A common objection could be that while meanings can in general be discriminated more or less finely, their individuation in terms of their logical properties (with logical equivalence as the only identity criterion) is unique. This objection assumes that there is something like a "logical degree of discrimination," but as soon as we give up truth-conditions there are only logical and extra-logical differences in discrimination; the former due to differences in the logical vocabulary and the latter due to differences in the extra-logical vocabulary.

four-valued logic discriminates more finely between propositions than classical logic does.

The above considerations can be used to formulate two independent objections against the suggestion that a focus on how finely propositions are carved out can be used to explain the specificity of an informational semantics. The first objection is that there is nothing specifically informational about logical discrimination. The second objection (raised in discussion by Graham Priest) is that if informational semantics is primarily about logical discrimination and how this affects the question of whether or not the proposition picked out by the conclusion is contained in the proposition picked out by the premises, then we should locate informational semantics at the level of algebraic semantics because in that framework the number of elements of the algebra determine the degree of logical discrimination in a more direct fashion.

The first objection is easily met by pointing out that since truth-conditional semantics require a non-classical account of truth (but see below for a further qualification of this claim), information-conditional semantics provide an account of logical discrimination that is more generally applicable. The second objection, by contrast, calls for a separate reply. The main worry in that case is that if algebraic semantics indeed provide the right level of analysis to formalise the content-nonexpansion platitude, all the talk about combining information from different situations might be entirely superfluous. As a result, there would be no need for a basic notion of objective information in an environment, and thus no need for an alternative reading of $s \Vdash A$. This would, in particular, cut the ties with one of the origins of information-conditional semantics, namely the interpretation of the ternary relation of Routley-Meyer semantics.

The claim that the connection between logical discrimination and deductive strength is central to algebraic semantics is correct and should not be denied. When used as an objection about where information-conditional semantics should be located, it is nevertheless misguided. Two complementary considerations explain why. On the one hand, we should stress that while logical discrimination is central to informational semantics, the resulting picture is only partial. Algebraic semantics are most useful to formalise this one aspect of how we access and use information, but this isn't the whole story. Because an algebraic semantics treats the relevant propositions directly, an informational interpretation of these propositions can only supply highly abstract pieces of information. As far as the algebraic semantics is concerned, pieces of information could as well be free-standing abstract objects like Fregean senses. The informational interpretation of the Routley-Meyer semantics rules out the latter interpretation by specifying that pieces of information find their origin in situations. This leaves the original algebraic semantics unaffected, but supplements it with an interpretation of the propositions as sets of situations. In summary, the informational interpretation of the Routley-Meyer semantics does not merely ensure that the abstract entities of the algebra correspond to (or are abstracted from) something real, but it also accounts for the fact that information arises in a context.

On the other hand, features like logical discrimination will surface in every formalism we use to study logical consequence; the only thing an algebraic

semantics does, is to make this very explicit. Indeed, we can think of an algebra as an abstraction of how finely propositions can be individuated in a Kripke-style semantics, and thus considerations about logical discrimination can immediately be applied to these kinds of structures as well. If we think of logical discrimination as an abstraction of how we can access and use information, then algebraic as well as relational (or Kripke) structures provide a formal model of this. Algebraic semantics do so in a very direct manner (all there is to the model are the elements, their ordering, and the choice of designated elements),[13] whereas relational structures do so in a manner where logical discrimination can be cashed out indirectly. This can be achieved by referring to different perspectives or different ways of accessing information (as in the previous section), or by focusing on global constraints on a class of models (as explained below).

As should be clear from the reply to one of the objections, we believe that the inverse relationship between deductive strength and logical discrimination makes more sense of an information-conditional semantics than of a truth-conditional semantics. This has to do with a difference in the order of explanation that is similar to the one described in the previous section, and with its use in an argument for logical pluralism.

To explain how the order of explanation can be reversed, we need to consider the notion of a global constraint imposed on a class of models. Let's start with a simple example:

i's being red implies that *i* is coloured.

This is a global constraint imposed by our language on the models that are appropriate for that language. It is a way to fix how finely propositions expressed by sentences using the words "red" and "coloured" are individuated. Given this global constraint, we cannot discriminate between situations where an object is merely red and situations where an object is red as well as coloured. It means that all situations are closed under this constraint. On the orthodox view, it is the logical space (e.g. the set of situations) which fixes the global constraints. Here, however, the global constraints come first and the structure of the logical space is built in such a way that it enforces these global constraints. According to this picture logical discrimination is conceptually prior to deductive strength.

The central role of logical discrimination (approached either in terms of perspective—access and use—or by looking at global constraints) is what distinguishes information-conditions from its main contenders. If we can make sense of the idea that information is a relational concept that depends on our discriminatory powers, and see logical discrimination as an abstraction of our ability to access and use information, we can claim that the extension of validity is a function of logical discrimination.

For a truth-conditional semantics, the main issue is that a difference in logic cannot reduce to a *mere* difference in logical discrimination. Since any difference in

[13] Remark that since by merely varying the set of designated elements, one can obtain different logics (compare for instance the strong Kleene 3-valued logic with the paraconsistent logic **LP**), the number of elements and their ordering do not suffice to fix the degree of logical discrimination.

logical discrimination needs a different account of truth, merely conceding that one can coherently disagree about the right or most appropriate way to individuate propositional contents doesn't suffice for logical pluralism. One must in addition concede that a coherent disagreement about truth is possible as well. Whenever a change in logic forces the adoption of a different theory of truth, being conservative about one's theory of truth strongly constraints one's logical options. By contrast, on a purely informational account, the choice of a theory of truth radically underdetermines the choice of a logic.[14]

The issue we want to raise here is best not confused with the related objection against Beall and Restall's pluralism that pluralism about logic forces one to be a pluralist about truth as well (Beall and Restall 2006, 100–2). To begin with, Beall and Restall are committed to there being different bodies of logical truths. But so are we, so this cannot distinguish our respective roads to logical pluralism. Instead, our point is that if one defines consequence in terms of truth-preservation over all *cases*, one has to accept that the ability to discriminate between two tautologies, or between two contradictions, does not only force one to accept a plurality of cases, but also that these cases are "*things* in which claims may be true" (Beall and Restall 2006, 89). Put in our terminology, this implies recognising that truth can be distributed, whereas we are only committed to the view that information is a distributed commodity. Similarly, recognising that deductive inference can fail to be cumulative one has to recognise that the premises we use in deduction no longer need to be true once used, whereas we are only committed to the view that our *access* to the information we use as a premise isn't always persistent. Perhaps Beall and Restall only hold that points of evaluations 'shift' when, for instance, we evaluate the truth of a relevant conditional, but do not take this to imply that premises used in deduction stop being true. If that's their position, then the connection between 'truth in a case' and 'truth simpliciter' is quite similar to our connection between truth and accessible information.

5 A Genuine Alternative!

What we've tried to establish is first and foremost that information-conditional semantics form a genuine alternative to truth-conditional and inferential semantics. By this we mean that despite sharing the formal machinery with truth-conditional semantics, namely a model-theoretic characterisation of consequence, information-conditional semantics isn't merely a more convenient variation of the standard truth-based reading of formal expressions like $s \Vdash A$.

The main obstacle to seeing our proposal as something more than a mere variant isn't the correct observation that an information-conditional semantics doesn't have a formalism of its own (even more, no account of logical consequence needs to be tied to a single dedicated formalism), but rather the additional assumption that the

[14] This is compatible with the fact that some theories of truth will turn a non-classical account of consequence into a non-classical account of truth (but see Mares (2008, sect. 10) on how we may add a truth-predicate), and with the fact that semantic paradoxes can still exclude some combinations of logics and formal theories of truth.

truth-conditional reading is the intended interpretation of any model-theoretic characterisation of consequence. Despite being the orthodox view, there's no reason why that reading should always be privileged. In our view, the inherent classicality of the truth-based reading of the model-theory merely reveals a question begging bias against non-classical logic. Truth-conditions are particularly compelling to the classical monist, but much less so for the pluralist.

We do not only believe that informational semantics is a viable alternative, but endorse the stronger view that as a philosophical semantics, information-conditions are the way to go for non-classical logicians and logical pluralists. One of its theoretical virtue lies in how it conceives of the disagreement between classical and non-classical logics. On a truth-conditional reading, the disagreement extends to a disagreement about the nature of truth. This often collapses paraconsistency with dialetheism, and pluralism about validity with relativism about the things that can make claims true or false. On a mixed reading (truth-based for classical logic, and information-conditional for non-classical logics), the disagreement concerns the subject-matter of logic. Finally, on an informational reading the disagreement can be considered as purely normative. It is about which logic is appropriate for the analysis of inferences. The benefit of recasting disagreements about the correct logic in these terms transcends the obvious advantages of avoiding the dialetheic peril or being able to advocate a non-relativist pluralist position.[15] It also yields an attractive and metaphysically unencumbered account of logical rivalry.

6 More Objections

Once it is established that an information-conditional semantics cannot easily be dismissed as a marginal improvement or variation of truth-conditional semantics, more substantial objections can be discussed. We conclude this paper by reviewing three types of objection.

A first objection is related to the scope of applicability of informational semantics, and raises the question whether an informational semantics does not misconstrue the relation between classical and intuitionist logic. This objection comes in two parts. The first part contends that because intuitionists really think the disagreement is metaphysical, it is a mistake to deflate this disagreement. The second part directly attacks the acceptability of an information-conditional semantics in virtue of its use of a model-theoretic characterisation of intuitionist logic (which we only mentioned in passing) that might not be acceptable to intuitionists. Hence, the outcome is misleading and this is due to the philosophical application of a formalism that doesn't provide a real insight into intuitionist logic (but see Aberdein and Read (2009) for a more extensive treatment of the different attitudes one may adopt). We recognise the philosophical difference between relevant and intuitionist logic, and do not want to obliterate the substantive

[15] Strictly speaking, if we follow the characterisation of Cook (2010), pluralism still implies a form of relativism in the sense that logical consequence is *relative* to a prior choice of a degree of logical discrimination.

disagreement between intuitionist and classical logic. We do however believe that, when complemented with a suitably neutral interpretation, model-theoretical formalisms can often be used for philosophical purposes. Informational semantics provide, especially if we keep in mind that it does not come with a dedicated formalism, a neutral interpretation in which it makes sense to explain intuitionist logic relative to a verificationist type of information. The resulting deflationary attitude towards metaphysical disagreement is intended, and entirely in line with the underlying motivations for logical pluralism.

A second objection calls the benefits of the informational conception of logic into question, and does so relative to, on the one hand, an inferential conception of consequence, and, on the other hand, Read's *homophonic* semantics for relevant logic (Read 1988). When compared to an inferential semantics, one indeed gets the impression that an information-conditional approach doesn't really yield a more attractive account of consequence. After all, the whole field of substructural logics—of which relevant logics are a specific type—was developed from a proof-theoretical perspective. This impression is misleading because it tacitly identifies the virtues of a proof-theoretical presentation with the virtues of an inferential account of consequence, and because it overestimates the importance of the model-theoretical formalism we used to capture the content non-expansion platitude. While we agree that someone who favours proof-theoretical tools might dislike our codification of the content-nonexpansion platitude, we also believe that the flexibility of the proof-theoretical formal framework to accommodate several non-classical logics does not directly carry over to the inferential conception of logical consequence. Indeed, whereas a truth-conditional account of relevant and intuitionist logic requires one to stretch the notion of truth one is using, an inferential account of, especially, relevant logics might require one to stretch the notion of proof one is using. For instance, while a proof-based reading of intuitionistically valid sequents is entirely natural, and perhaps extends to several other logics, an informational reading might be more natural for relevantly valid sequents. To a certain extent the issue is analogous to what we already pointed out with respect to truth-conditional accounts. On the orthodox account where proof-theory formalises a proof-conditional account of consequence, one can only combine relevant consequence with an inferential conception of consequence by stretching one's conception of proof to account for modes of inference that aren't cumulative. This can be avoided by replacing the common proof-based reading of sequents by an informational reading.[16]

When compared to Read's analysis of logical consequence as the impossibility for the fusion of the premises to be true *and* (read as "fusion") the conclusion to be false, the benefits of the informational conception are not so clear. Both use the fusion-connective, and are thus similar in that respect. Their mutual disagreement seems to boil down to the following. According to Read's *Scottish Plan* for relevant logic (Read 1988), we have to give up on an extensional meta-theory to keep an

[16] Paoli (2002, 3.3) proposes such a reading, but because he uses the term 'information' to refer to data-types, it is only appropriate in the absence of the structural rules of weakening and contraction. A more open-ended reading is required.

account of relevant consequence that is based on truth-preservation. By contrast, on the informational account an extensional meta-theory can be retained independently of the logic one chooses, but the truth-preservation platitude has to be given up. On the former account, the frame semantics devised by Routley and Meyer merely serve a technical purpose. This creates a gap between the tools that are used to prove results about a logic, and the methods that are used to give a philosophical account of the consequence relation. One benefit of the adoption of an informational semantics is that there is no such a gap (though we agree that if one already believes that the Routley-Meyer semantics is a gentile semantics, it is only natural that there should be such a gap).

A third objection (adapted from a comment made in discussion by Agustín Rayo) targets the presumed need for a primitive notion of information. Specifically, the objection contends that it isn't quite clear what such a primitive notion would add to a picture where Stalnaker-propositions (sets of possible worlds) provide the foundation for objective information and more fine-grained ways of carving out propositions are obtained by relativizing attributions of Stalnaker-propositions to the purposes of a particular project (the "issues under consideration", see Perry (1986)) rather than by supplying additional structure of the kind that allows situations to be included in worlds. In our terminology: Contexts may change and this affects the distinctions that can be made, but these are changes that do not correspond to a difference in perspective. This, for sure, would make a more primitive notion superfluous, and result in an even simpler informational semantics. From a formal point of view, there is perhaps not much to object to this proposal. Even the objection that this would equivocate between information and informational content could be avoided by pointing out that there is nothing more to information than the individuation of a set of possible worlds, and that the only difference between information and informational content is just the difference between a qualitative individuation and a (quantitative) measure of what is individuated. But even then there are several things to worry about.

One such worry is that on this type of account the world, as it were, already privileges a way of carving out propositions. If the world already picks out a certain level of abstraction,[17] the fine-structure we add to this picture can only be due to our own failure to recognise that certain propositions are really *the same*. This idea is not only opposed to the proposal that information is a relational notion, but it also reduces any further distinction to something we impose on the world. Such distinctions all too easily become distinctions out of ignorance (or at least distinctions that are exclusively prompted by our cognitive abilities and purposes) rather than distinctions that become available because of what the world is like and how we interact with it. As a consequence, we would have to give a very one-sided

[17] As Stalnaker puts it: "The formalism of possible worlds semantics assumes that possible states of the world are disjoint alternatives, and that everything that can be said within a given context can be said by distinguishing between these alternatives. This assumption of internal completeness is required by the explanation of propositional contents as sets of possible states of the world, and this explanation is motivated by our account of the nature of representation: since to represent the world just is to locate it in a space of alternative possibilities, content should be explained in terms of those possibilities." (Stalnaker 1986, 118).

explanation of why making additional distinctions can sometimes give us a better grip on the world. Having a primitive notion of information that leaves open the right level of abstraction does not have such problems.

A second worry—though one that many would consider a positive outcome of the theory—is that on such an account classical logic would indeed have a special status. This would threaten the prospects for a logical pluralism based on the informational conception of logical consequence, and would again yield a rather one-sided account of non-classical logics. Specifically, the need to balance logical discrimination and deductive strength could not be explained by the respective epistemic virtues of having strong theories or having very refined theories, but would reduce to having to bring our logic in line with our epistemic abilities and cognitive enterprises.

Acknowledgments This paper was first presented at the *Conference on the Foundations of Logical Consequence* that was held in 2010 at the University of St Andrews. We would like to thank the organisers and audience of this conference as well as two referees of this journal for their encouragement and valuable feedback.

References

Aberdein, A., & Read, S. (2009). The philosophy of alternative logics. In L. Haaparanta (Eds.), *The development of modern logic* (pp. 613–723). Oxford: Oxford University Press.

Allo, P. (2007). Logical pluralism and semantic information. *Journal of Philosophical Logic, 36*(6), 659–694.

Anderson, A. R., Belnap, N. D., Dunn, J. M. (1992). *Entailment. The logic of relevance and necessity* (Vol. II). Princeton: Princeton University Press.

Barwise, J. (1993). Constraints, channels, and the flow of information. In P. Aczel, D. Israel, S. Peters, Y. Katagiri (Eds.), *Situation theory and its applications* (Vol. 3, pp. 3–27). Stanford: CSLI.

Barwise, J. (1997). Information and impossibilities. *Notre Dame Journal of Formal Logic, 38*(4), 488–515.

Barwise, J., & Perry, J. (1999). *Situation and attitudes*. Stanford: CSLI Publications.

Beall, J. C., & Restall, G. (2006). *Logical pluralism*. Oxford: Oxford University Press.

Beall, J. C., Brady, R., Dunn, J. M., Hazen, A., Mares, E., Meyer, R. K., et al. (2011). On the ternary relation and conditionality. *Journal of Philosophical Logic*, Online First. doi: 10.1007/s10992-011-9191-5.

Belnap, N. (1976). How a computer should think. *Oxford international symposium on contemporary aspects of philosophy*. Oxford.

Carnap, R., & Bar-Hillel Y. (1952). *An outline of a theory of semantic information*. MIT, p 247.

Cook, R. T. (2010). Let a thousand flowers bloom: A tour of logical pluralism. *Philosophy Compass, 5*(6), 492–504.

Copeland, B. J. (1979). On when a semantics is not a semantics: Some reasons for disliking the Routley-Meyer semantics for relevance logic. *Journal of Philosophical Logic, 8*(1), 319–413.

Dretske, F. (1999). *Knowledge and the flow of information*. Stanford: CSLI.

Dunn, J. M. (1976). Intuitive semantics for first-degree entailments and 'coupled trees'. *Philosophical Studies, 29*, 149–168.

Dunn, J. M. (1993). Star and perp: Two treatments of negation. *Philosophical Perspectives, 7*, 331–357.

Floridi, L. (2005). Is information meaningful data? *Philosophy and Phenomenological Research, 70*(2), 351–370.

Goble, L. (2003). Neighborhoods for entailment. *Journal of Philosophical Logic, 32*(5), 483–529.

Hanson, W. H. (1980). First-degree entailments and information. *Notre Dame Journal of Formal Logic, 21*(4), 659–671.

Humberstone, I. L. (2005). Logical discrimination. In J.-Y. Béziau (Eds.), *Logica Universalis* (pp. 207–228). Basel: Birkhäuser Verlag.

Israel, D., & Perry, J. (1990). What is information. In P. Hanson (Eds.), *Information, language and cognition* (pp. 1–19). Vancouver: University of British Columbia Press.
Kemeny, J. G. (1953). A logical measure function. *The Journal of Symbolic Logic, 18*(4), 289–308.
MacKay, D. M. (1969). *Information, mechanism and meaning*. Cambridge, MA: MIT Press.
Mares, E. D. (1997). Relevant logic and the theory of information. *Synthese, 109*(3), 345–360.
Mares, E. D. (2004). *Relevant logic. A philosophical interpretation*. Cambridge: Cambridge University Press.
Mares, E. D. (2008). *Information, negation, and paraconsistency*. Melbourne: WCP4.
Mares, E. D. (2009a). General information in relevant logic. *Synthese, 167*(2), 343–362.
Mares, E. D. (2009b). Relevance and conjunction. *Journal of Logic and Computation*. doi: 10.1093/logcom/exp068.
Mares, E. D. (2010). The nature of information: A relevant approach. *Synthese, 175*(Sup. 1), 111–132.
Paoli, F. (2002). *Substructural logics a primer*. Dordrecht: Kluwer.
Paoli, F. (2007). Implicational paradoxes and the meaning of logical constants. *Australasian Journal of Philosophy, 85*(4), 553–579.
Perry, J. (1986). From worlds to situations. *Journal of Philosophical Logic, 15*(1), 83–107.
Priest, G. (2001). Logic: One or many? In J. Woods, & B. Brown (Eds.), *Logical consequence: Rival approaches* (pp. 23–38). Stanmore: Hermes.
Read, S. (1988). *Relevant logic. A philosophical examination of inference*. Oxford: Basil Blackwell.
Read, S. (2003). Logical consequence as truth preservation. *Logique & Analyse, 46*(183–184), 479–493.
Restall, G. (1995). Information flow and relevant logic. *Logic, Language and Computation: The 1994 Moraga Proceedings*.
Restall, G. (1996). *Notes on situation theory and channel theory*.
Restall, G. (2005). Logics, situations and channels. *Journal of Cognitive Science, 6*, 125–150.
Saguillo, J. M. (2009). Methodological practice and complementary concepts of logical consequence: Tarski's model-theoretic consequence and Corcoran's information-theoretic consequence. *History and Philosophy of Logic, 30*(1), 21–48.
Stalnaker, R. (1986). Possible worlds and situations. *Journal of Philosophical Logic, 15*(1), 109–123.
Turvey, M. T., & Carello, C. (1985). The equation of information and meaning from the perspectives of situation semantics and Gibson's ecological realism. *Linguistics and Philosophy, 8*(1), 81–90.
Wansing, H. (1993). Informational interpretation of substructural logics. *Journal of Logic, Language and Information, 2*, 285–308.
Wansing, H. (1993). *The logic of information structures*. Berlin: Springer.