Informational content and information structures: a pluralist approach

Patrick Allo
CLWF – VRIJE UNIVERSITEIT BRUSSEL
IEG – OXFORD UNIVERSITY
patrick.allo@vub.ac.be

Abstract

The tutorial connects two notions of information: the inverse relationship principle which relates informational content to the exclusion of possibilities, and information-structures based on a partial ordering on states of information. Jointly, these allow the formulation of several distinct precise notions of content-individuation.

1 Introduction

There are several ways in which an agent can interact with information. Not only can an agent have information, information can also be received, passed-on, combined with other pieces of information, etc. Additionally, there is a sense in which information can be valuable, or fail to be valuable. Its receipt can be useless because it was already received at a previous occasion, or even be useless simply because it was impossible for that agent not to hold that information. In both these cases, information (or at least its receipt) can be called redundant. Finally, there is also a sense in which information may simply be irrelevant. Even though not useless in any absolute kind of way, it might just not fit an agent's interests at a certain time.

Disregarding the latter aspect of relevance, we can disambiguate between three types of relations which describe conceptually distinct ways an agent relates to, and interacts with, pieces of information (Floridi (2006)).

Being Informed The state in which an agent holds the information that p.

Becoming Informed The action or process of receiving the information that p.

Being Informative The assessment that the information that p has non-zero content.

Floridi describes the distinction between the first two as the distinction between the *statal* condition of holding information, and the *actional* process of receiving information (Floridi (2006): sect. 1). The connection between both is rather obvious as becoming informed that p represents the action which triggers the change between possibly not being informed, and (if successful) necessarily being informed that p.

Yet, both in everyday parlance and in more systematic treatments of information, these statal and actional interpretations are usually cashed out as the distinction between information and knowledge. Recently, with the development of dynamic epistemic logics, this connection between a dynamic notion of information and the static notion of knowledge has become a central concern within the logical community. Especially with regard to the basic cases where new information is objective or purely factual,[1] it is fair to say that the interaction between the static and the dynamic aspects of knowledge, belief and information is by now formally well-understood (Baltag & Moss (2004), and the tutorial by B. Kooi).

At first sight, the notion of being informative seems quite distinct from the latter two. It is neither a specific state nor an action that relates to such states. As such, it is more often described using terms like accuracy, reliability or value than by means of the relations of holding and receiving information. In the end, this last notion seems hardly concerned with the information itself and how communicators interact with it. Instead it aims at the bare question of *how much information* there is. Traditionally the property of being informative is akin to the measuring (and individuating) of informational content, and its investigation has lead to the development of several quantitative theories of informational content. The best-known representative of the latter approach is Carnap & Bar-Hillel's *Theory of Semantic Information* (Carnap & Bar-Hillel (1952)) which was developed as a semantically sensitive modification of Shannon & Weaver's (1998) *Mathematical Theory of Communication*.

A basic idea of what approaches in the line of Carnap & Bar-Hillel's work aim at, is obtained when we focus on what they have in common; namely their explication of content as the potential to make selections. This idea finds its most general formulation in the *Inverse Relationship Principle* which relates information to alternatives or possibilities, and is the basic tenet of what we have come to know as the *information as range* paradigm (Van Benthem (2006+)).

> Whenever there is an increase in available information there is a decrease in possibilities and vice versa. (Barwise (1997): 491)

What is characteristic of this description, is that there is no need to explain informational content relative to a previously given set of alternatives a receiver might entertain, for it can straightforwardly be made depend on *all*

possible alternatives. Obviously, if what matters is the potential to exclude rather than what is actually excluded, one cannot but take all alternatives into account. Perhaps this feature of how exactly information is selective provides the best illustration of the thriving idea behind the notion of informational content: it primarily treats content as a property of sentences or propositions. No reference to an actual informee needs to be made, and this makes the notion of informativeness (or being informative) different from the statal and actional notions of being and becoming informed.

Of course, this should not entail that this very property of propositions is irrelevant to how cognitive agents interact with pieces of semantic information. To be precise, informativeness and specifically the property of being informative can easily be related to the property of being a non-vacuous update. Indeed, if i stands for an agent's state of information, and $i \uparrow A$ for its resulting state after being updated with the (for this purpose, purely propositional) information that A, A is informative iff the latter update is non-vacuous; that is, if $i \neq i \uparrow A$. Consequently, if the latter holds at least for some agent, it can be concluded that A is informative *simpliciter*—it has non-zero informational content. If one considers the formal details, this is all but surprising; the notions of content and update conceive of information in exactly the same 'possibility-excluding' sense. Even so, it remains an important insight that the problem of individuating the informational content conveyed by a piece of *Declarative, Objective and Semantic Information* can equally well be tackled in terms of updates. This idea fully agrees with the now popular view that "information is what it does."

The view defended in this paper is that as an explanation of informativeness, the *Inverse Relationship Principle* is somehow incomplete. Putting the dogmatic adherence to its classical precisification aside, the intuitively appealing connection between information and the exclusion of alternatives is demonstrably insufficient for the precise (and determinate) individuation of the content of a given piece of information. The cause for this under-determination of content-individuations by the *Inverse Relationship Principle* can be retraced to the different ways in which a given space of alternatives (the agent's modal epistemic space) can be represented. On the pure formal level, this theoretical plurality can be related to how generic information-structures generalise the classical possible worlds approach to epistemic possibility. To turn this mere theoretical plurality into a fruitful explication of informational content, the actual modal space we use to model an agent's state of information is related to the Level of Abstraction (henceforth, LoA) we adopt to model a given system. Intuitively, this notion of a LoA is best considered as, in the first place, a function of (the modeller's) ability to tell two states of information apart, and, in a second

time, as a function of an agent's presumed ability to discriminate between the content of two pieces of information.

2 Information and the exclusion of alternatives

As we wish to reveal the previously mentioned incompleteness that arises from the relation between information and the exclusion of alternatives, it is necessary to adopt a sufficiently unbiased formulation of the principle which expresses that relation. The *Generalised Inverse Relationship Principle* (henceforth, *GIRP*) introduced below captures such a basic and unassuming reading of the initial insight.

(GIRP) The informational content of a piece of information is given by the set of cases it excludes.

Even though this version should be intuitively acceptable, we might still point out how it avoids the implicit presuppositions of some of its main contenders. The benefits mainly revolve around (i) the reference this principle makes to *cases* rather than to *alternatives* or *possibilities*, and (ii) the use of the term *exclusion* rather than *inconsistency*. As for the former, it should be clear that—partly due to its own vagueness—the notion of a *case* does not not share the need for exclusiveness and exhaustiveness we often associate with the term *alternative*, nor does it privilege the completeness and consistency suggested by the term *possibility*. For the latter, then, the notion of exclusion turns out to be equally neutral towards the many issues that involve the properties of negations. Since only for the standard version we shall elucidate in the next section it holds that the cases excluded by a piece of information are exactly the cases which are inconsistent with or support the negation of that piece of information, this kind of generality is clearly beneficial. As a consequence, and this might sometimes be overlooked, the use of the term exclusion is required to leave open the precise interpretation of the term *case*.[2]

The *Generalised Inverse Relationship Principle* we just introduced can easily be related to Beall & Restall's *Generalised Tarski Thesis*, for their argument for a pluralist understanding of logic equally depends on the alleged incompleteness of an intuitively valid principle. This moreover entails that a robust case for the incompleteness of *GIRP* could in principle be entirely motivated as a mere corollary of Beall & Restall's logical pluralism. In order to construct such a case, the move from logical pluralism to the conclusion that the *GIRP* does not uniquely determine the content of a piece of information ought to be made dependent on the assumption that tautologies (and logical truths in general) are devoid of any informational content.

This very old idea can at least be retraced to the Traktarian affirmation of the "emptiness of logic."

Yet, this is not to say that a more direct argument cannot be made. Rather than depending on the presumed uninformativeness of logic,[3] the remark that information is always assessed at a given level of abstraction can be used to arrive at the same (and perhaps even a more general) conclusion. For now, all we need to say is that by avoiding the association with a pluralism about logical consequence, several references to controversial issues in the philosophy of logic can be avoided or at least bypassed (most importantly, the precise interpretation of the platitude that, at least in its deductive sense, logic is informationally empty or devoid of any factual content). At least for argumentative reasons, it proves much less problematic to relativize content-assessments, than to relativize the norms which govern valid deductive inference. One reason for this, is that whereas content-assessments primarily serve a theoretical purpose, our theorizing about validity is more often related to practical concerns like the norms for deductive reasoning in the vernacular. Put differently, although both validity and informational or semantic content are technical terms, the latter does not come with many intuitions that go beyond the inverse relationship principle itself. By contrast, logic traditionally comes with something like an intended interpretation and a canonical domain of application. This fact makes it particularly hard to resist the view that either the conclusion follows from the premises, or it does not follow from it (see e.g. Priest (2001); Read (2006)). By avoiding similar constraints, a more versatile pluralism can be upheld. For now, it suffices to say that such a pluralism is an instance of the "pluralism without relativism" we can derive from the method of abstraction (Floridi, L. and J. W. Sanders (2004)).

Despite the above suggestion that there is no apriori reason to single out the classical interpretation of the *Inverse Relationship Principle* as its one true formalisation, its overall simplicity is a good reason for adopting it as the starting point of our present investigation. In its original version, the classical approach is due to Carnap & Bar-Hillel (1952) and makes use of so-called *state descriptions* to formalise the cases or alternatives the *GIRP* refers to. The latter are conjunctions of literals (atomic propositions or their negation) that describe exactly one possible state of affairs. In other words, these serve as the syntactic descriptions of a *case*, and specifically as a description of a maximally consistent one. Rephrased in semantical terms, the classical approach interprets the *cases* referred to in the *GIRP* as possible worlds or as classical models. Of these, the version which refers to models (see Kemeny (1953)) can serve as a more versatile and elegant version of Carnap

& Bar-Hillel (1952)'s original proposal.

In this presentation, we restrict ourselves to its qualitative version: the individuation of a proportion of the total modal space rather than the numerical quantification of informational content. Let A be a piece of information expressed in a standard propositional language, then its informational content is given as:

DEFINITION 2.1 (CLASSICAL CONTENT) $\mathrm{CONT}(A)$ = the set of models in S (the set of all models) inconsistent with A, namely $\{s \in S \mid s \Vdash\!\!\sim A\}$ or equivalently $\{s \in S \mid s \nVdash A\}$. Consequently, we obtain:

$$\begin{aligned} \mathrm{CONT}(\top) &= \emptyset, \text{ for } \top \text{ is a classical tautology.} \\ \mathrm{CONT}(\bot) &= S, \text{ for } \bot \text{ is a classical contradiction.} \end{aligned}$$

As previously suggested, a similar result is obtained with updates.

DEFINITION 2.2 (CLASSICAL UPDATES) The content of A is the set of cases ruled out by updating the set of all cases with A, namely: $S \setminus (S \uparrow A)$. Consequently, we obtain:

$$\begin{aligned} \mathrm{CONT}(A) = \emptyset \quad &\text{iff} \quad (S \uparrow A) = S \\ \mathrm{CONT}(A) = S \quad &\text{iff} \quad (S \uparrow A) = \emptyset \end{aligned}$$

Usually, the standard objection to this classical approach refers to its plainly unintuitive identification of any contradictory piece of information with the maximally informative piece of information (to which Floridi (2004) refers as the *Carnap Bar-Hillel Paradox*). This line of reasoning largely parallels one of the best known objections to classical logic; viz. the rejection of the classically valid implication of all propositions by a contradiction (*ex falso quodlibet*). Additionally, since the dual of this objectionable property (the identification of any tautological piece of information with the minimally informative piece of information) can similarly be related to some of the so-called fallacies of relevance classical logic has repeatedly been blamed for, the analogy with standard arguments for logical revision proves rather strong. Promising as this approach may seem, it can only be turned into a rejection of the classical approach to content by making an appeal to a substantive connection between logical consequence and informational content. And even then, it would only motivate a revision of the classical account, not a pluralist interpretation. An argument that proceeds independently is therefore preferable.

The ground for a more general objection to this classical proposal is already implicit in the formulation of the discontent with how this approach handles contradictory pieces of information. That is, the content

of every piece of contradictory information is identified with *the same* fragment of the relevant space of possibilities: namely the totality of that space. Such a failure to discriminate between the content of two apparently distinct pieces of information (say, an overt contradiction and the negation of a fairly complicated logical truth) is reminiscent of older issues in the philosophy of logic and in formal semantics: the problems of hyperintensionality, granularity, and more generally the fact that meanings cannot straightforwardly be identified with intensions. Equally often, similar issues are related to the problem of logical omniscience in modal epistemic logic. In short—and in a highly naive formulation—the relevant problem thus surfaces as the inability of a modelling framework to be as discerning or discriminating as that fragment of the real world it intends to model. A solution to this problem is often thought to depend on the correct evaluation of the trade-off between accounting for each syntactical difference and adopting an overly crude semantical approach which disregards any syntactical particularity. In many cases this view implies that some determinate position between these two extremes suffices to account for most, if not all, relevant linguistic phenomena.

As a discussion of the drawbacks of the standard approach to informational content, the exposition in Sequoiah-Grayson (2006) is especially attentive to this more general issue of how one should evaluate the information yield of two logically equivalent pieces of information. Again, this concern can quite easily be reduced to the traditional debates on the presumed analyticity of logic and the relation between logical equivalence and synonymy, but perhaps it makes more sense to situate it in the broader context of finding the correct level of abstraction for the assessment of the content conveyed by a piece of information. Crucially, the introduction of the method of abstraction as a basic methodological lead helps us to radically transform the old problem of finding the 'right level' of abstraction needed for semantic theorising. By shifting the focus from the defence of a certain level of abstraction to, on the one hand, the explicitation of its ontological commitments, and, on the other hand, how one such level coheres with other levels of abstraction, the question of what counts as the right level is no longer on the foreground. This approach suggests more than a mere methodological pluralism. To describe the impact of its adoption with a slogan, it could be said that it favours a pluralist approach, but steers clear of a relativist understanding of informational content. Concretely, this means that while the semantic content conveyed by a piece of information is to a certain extent independent of the actual informees that are involved, reality does not (and cannot) enforce a single particular way of carving up the space of possibilities relative to which the content of that piece of informa-

tion is evaluated. Despite the overt attention to levels of abstraction, such pluralist considerations play no role of importance in mainstream semantical theories. Several well-known examples, as the one sketched below, illustrate this.

Traditionally, the apparent failure of the principle of substitutivity for propositions within the scope of an attitude verb has served as the most urgent example which motivated the search for a sufficiently fine-grained semantical framework (think of the puzzles due to Frege and Kripke). The problem usually surfaces in the following form. Where A stands for some attitude verb, given the assertion that $\mathsf{A}_a A$, one is not always prepared to assert that $\mathsf{A}_a B$ in virtue of the fact that A and B are logically (or materially) equivalent. In more explicit terms, if a certain context is modelled using a possible worlds framework, for all A and B true at exactly the same possible worlds, that model cannot but treat the attitude reports that $\mathsf{A}_a A$ and that $\mathsf{A}_a B$ as being essentially the *same* attitude report or at least reports of two identical attitudes. Precisely this conclusion is often rejected because of its failure to recognise the fact that within such a context we are effectively confronted with one attitude report being warranted while the assertion of a logically equivalent one is unacceptable or at least awkward.

With this example in mind, and despite the methodological maxims inspired by the method of abstraction, it is obvious that one cannot appeal to the sheer theoretical feasibility and consistency of several ways to carve up a given space of possibilities to dismiss the question of which level of abstraction one should adopt. To properly appreciate why the problem of fixing the proper level of abstraction does not immediately surface with regard to the evaluation of informational content, we should contrast the problem as it arises with the attitude reports with the more general problem of content-assessments. The point that needs to be highlighted bears on the following discrepancy. An attitude report is potentially both an assertion about a system or a context *and* an assertion that can be made within that context. Therefore the formal model we use to explain the semantics of, say, the verb 'know' should not only accurately reflect what it means for an agent to know, but should (pragmatic phenomena aside) be equally accurate with regard to the functioning of knowledge ascriptions made within the relevant system. Quite differently, content assessments typically do not have to function within the system they refer to. When this divergence between the double role of attitude reports and the exclusively system-external role of content assessments is duly appreciated, one can easily recognize why a theory of propositional attitudes cannot as easily be constructed at any conceivable level of abstraction. As a semantic

theory, there is a system-internal functioning which needs to be accounted for.

Two conclusions follow from this insight. First, one can suggest that as a formal model of system-external attitude reports, the possible worlds model cannot be blamed for being outrightly mistaken. It is just a model of knowledge at a high level of abstraction, and incidentally, that high level of abstraction does not seem to match the (implicit) level of abstraction at which system-internal knowledge ascriptions are made. Some might even claim that this is a problem for semanticists interested in the attitude verbs, but therefore not necessarily a problem for someone interested in knowledge itself. This attitude largely coincides with the usage of epistemic logic in computer science and economics. Such approaches are, rather than being outrightly mistaken, highly abstract or idealised. Moreover, they are so for several good reasons like simplicity or tractability.

Second, when it comes to content-assessments, there are no prior data which need to be accounted for. We surely have intuitions about how some pieces of information are more informative than others, but such intuitions are rarely grounded in or reflected by a distinct linguistic practice of making content-assessments. Even so, this does not mean that a theory of informational content is irrelevant to our understanding of a broader communicative practice (for this would mean that a general theory of informational content would have no purpose at all). All that needs to be emphasised is that the generic practice of evaluating the content of information conveyed by the agents of a certain system does not need to stand in any direct and transparant relation to the communicative practice of these agents. As a consequence, and notwithstanding the fact that this point should be further refined, it is highly implausible that a unique level of abstraction could be solely derived from the functioning of the system.

By way of conclusion, the benefits of conceiving a theory of informational content as a general theory of a modeller's content-assessments are, once more, contrasted with two better known theoretical enterprises. These are the study of logical consequence and the study of cognitive attitudes; two enterprises which are commonly (but not uncontroversially) thought to yield theories that are directly relevant to cognitive practice. As such, the specification of the semantics of attitude verbs is considered relevant to practical and theoretical reason in general, whereas the formal concept of logical consequence is considered normatively and descriptively relevant to the practice of deductive reasoning. On the assumption that the above connections are as tight as the received view has it, a solid case can be made for building these theories at a level of abstraction which closely matches the implicit level of abstraction of the system-internal practice. Omitting

the technicalities, this is just to say that our theories of cognitive attitudes or of logical consequence can only be sound and complete with regard to the practice they intend to explain provided they are formulated at a level of abstraction that is sufficiently close to the agent's actual ability to make distinctions. Referring to a suggestion made in the introduction, in these cases one is thus required to take the agent's cognitive abilities and cognitive restrictions as the ultimate guide for deciding what counts as the right level of abstraction for a model.[4]

In that perspective, and only when used in conjunction with the appropriate bridge principles, an abstract theory of logical consequence should only sanction a deductive inference if it is acceptable for an agent to reason accordingly. Likewise, from a modellers perspective we should in general not ascribe knowledge (or a suitably strong form of belief) to an agent if by making such an ascription we would falsely predict that agent to be warranted in making certain assertions or taking some action, or even predict that another agent would be warranted in ascribing knowledge to that first agent. Thus, even if the relation between practical and theoretical reason is a highly controversial issue, once we settle on their relation we cannot but require our theoretical reconstructions of deductive consequence and knowledge to stand in the proper relation to their practical siblings and the empirical evidence we have for that practice.

Admittedly, the absence of a well-defined practical activity whose functioning could be explained on the basis of content-assessments is not enough to prove the generality of content-assessments. As it suggests the conceivability of practical tests for informativeness, the tight connection between informativity and the distinction between vacuous and non-vacuous updates makes this concern even more explicit. However, since this is not the place for a full-blown defence of the generality or not *out-of-the-box* applicability of content-assessments, two elementary considerations should do. The first point is that reasoning about information is more intimately connected to te practice of modelling than, say, our reasoning about knowledge and belief is. This suggests that the modeller's use of informational concepts ranges well beyond the common usage of intentional idioms, and perhaps even ranges beyond the already fairly general usage of these idioms that is motivated by the adoption of Dennett's *intentional stance*. Secondly, it is important to realise that the choice of a certain level of abstraction is not only determined by the intrinsic properties of the system. The choices the modeller makes and the constraints under which she operates are indeed of equal importance. When we are in the modelling business instead of pursuing traditional conceptual analysis, both the adoption of a coarser model as well as its refinement are among the acceptable practices.

3 INFORMATION STRUCTURES

Although we argued that the GIRP captures one of our most pervasive intuitions regarding informativeness, it is by no means the only way to reason qualitatively about information. Perhaps, the fact that states of information can be ordered by a relation of *information inclusion* provides an even more abstract perspective on information and information increase in particular. Relying on such orderings, a modal space can be carved up in many more ways than the standard possible worlds picture suggests. These refined structures can subsequently be used in conjunction with the general insight expressed by the *inverse relationship principle*. To a first approximation, the following setup captures the basic concepts we need for that refinement. Where S is the set of all states of information, $\sqsubseteq$ defines a partial order on S that can be interpreted as a relation of information-inclusion if it satisfies a so-called monotonicity condition. Intuitively, $s_1 \sqsubseteq s_2$ expresses that s_2 contains at least as much information as s_1 does, and hence two states are identical iff $s_1 \sqsubseteq s_2$ and $s_2 \sqsubseteq s_1$. Since our usage of the method of abstraction not only relates to the ability to discriminate states of information, but also to the ability to discriminate actual pieces of information, we need to include a further identity-criterion to that effect. For that purpose, propositions are identified with a particular class of subsets of S, namely those that are upwardly closed under $\sqsubseteq$: $A \subseteq S$ is a proposition iff $s_1 \in A$ and $s_1 \sqsubseteq s_2$ jointly imply $s_2 \in A$. Thus, as one might expect, differences in informational content relate to the ability to tell two propositions apart which itself is a function of our identity-criterion for states of information.

While this approach can be exploited in several ways, we shall only focus on one such structure and its relation to the standard possible worlds model.

Our starting point is a Routley-Meyer Frame, a peculiar extension of the information-structures we previously introduced.

DEFINITION 3.1 (ROUTLEY-MEYER FRAME) A Routley-Meyer frame is a structure $(S, Log, \sqsubseteq_{\_}, *)$, where S is a set of situations, Log the subset of logical situations in S, $\sqsubseteq_{\_}$ a ternary relation on S,[5] and $*$ a unary operation on S. A partial order $\sqsubseteq$ on S is defined as $s_1 \sqsubseteq s_2 := \exists c \in Log \;\&\; s_1 \sqsubseteq_c s_2$, such that $(S, \sqsubseteq)$ is a bare information-structure which satisfies the elementary monotonicity-condition

$$s_1 \in p \;\&\; s_1 \sqsubseteq s_2 \quad \Rightarrow \quad s_2 \in p$$

This is already sufficient for defining the satisfaction conditions for atoms and the extensional conjunction and disjunction $\sqcap$ and $\sqcup$.[6] We say that an evaluation $\Vdash$ is a relation between the set of situations S, and the set of

propositional formulae.

$$\begin{array}{rcl} s \Vdash p & \text{iff} & s \in p \\ s \Vdash A \sqcap B & \text{iff} & s \Vdash A \text{ and } s \Vdash B \\ s \Vdash A \sqcup B & \text{iff} & s \Vdash A \text{ or } s \Vdash B \end{array}$$

So far, the hereditary condition remains satisfied. If the satisfaction conditions for negation are given by:

$$s \Vdash \sim A \quad \text{iff} \quad s^* \nVdash A$$

the hereditariness of negated formulae and double-negation elimination are respectively enforced by:

$$\begin{array}{rcl} s_1 \sqsubseteq s_2 & \Rightarrow & s_2^* \sqsubseteq s_1^* \\ s & = & s^{**} \end{array}$$

The (asymmetric, transitive and non-reflexive) ternary relation $\sqsubseteq_{\_}$ which is required for the satisfaction-conditions for the intensional connectives should in its turn satisfy the following constraints to ensure respectively the additional hereditariness conditions, assertion, contraposition, and interchange.

$$\begin{array}{rcl} s_1 \sqsubseteq_c s_2 \;\&\; c' \sqsubseteq c & \Rightarrow & s_1 \sqsubseteq_{c'} s_2 \\ s_1 \sqsubseteq_c s_2 & \Rightarrow & c \sqsubseteq_{s_1} s_2 \\ s_1 \sqsubseteq_c s_2 & \Rightarrow & s_2^* \sqsubseteq_c s_1^* \\ \exists s \in Log \;\&\; s_1 \sqsubseteq_{c_1} s \sqsubseteq_{c_2} s_2 & \Rightarrow & \exists c \in Log \;\&\; s_1 \sqsubseteq_c s_2 \;\&\; c_1 \sqsubseteq_{c_2} c \end{array}$$

The evaluation $\Vdash$ is then extended with the following clauses:

$$\begin{array}{rl} s \Vdash A \to B & \text{iff } \forall s_1, s_2 \in S \text{ where } s_1 \sqsubseteq_s s_2, \quad s_1 \Vdash A \Rightarrow s_2 \Vdash B \\ s \Vdash A \oplus B & \text{iff } \forall s_1, s_2 \in S \text{ where } s_1 \sqsubseteq_s s_2, \quad s_1^* \nVdash A \Rightarrow s_2 \Vdash B \\ s \Vdash A \otimes B & \text{iff } \exists s_1, s_2 \in S \text{ where } s_1 \sqsubseteq_{s_2} s, \quad s_1 \Vdash A \;\&\; s_2 \Vdash B \end{array}$$

Jointly, this yields the semantics for distributive sub-exponential linear logic or contraction-free relevant logic. In addition to its rejection of *ex falso quodlibet* and several other irrelevancies, the main benefit of this system is its enhanced degree of logical discrimination. The prime example of this enhanced discriminatory power is its ability to discriminate between the truth of 'A OR B' in virtue of the truth of either A or B (viz. $A \sqcup B$) and in virtue of the correctness of the inference that if NOT $-$ A then B (viz. $A \oplus B$). This fully agrees with the truism discussed by Humberstone (2005) that in general inferential strength and logical discrimination are inversely proportional.

With regard to the aim of defining alternative precisifications of the *GIRP*, Routley-Meyer frames have two major benefits. The first benefit derives from the fact that the logical points in S can be understood as *situations*, states of affairs which need not be maximally consistent, but instead can be both incomplete and inconsistent. Thus, it need not be true that either $s \Vdash p$ or $s \Vdash\!\sim p$. Dealing with partiality by means of logical points which can themselves be partial is particularly interesting, for it immediately agrees with the idea that when *information reports* express partial information, they do so relative to a situation, not to the world as a whole. This approach is essentially due to Barwise & Perry's situation semantics, and explicitly defended in Israel & Perry (1990). The connection with the semantics for relevant logic was already anticipated in Dunn (1976), and specifically related to the Routley-Meyer frames by Restall (1995) and Mares (1997). This benefit is exemplified by the notion of situated content that can be defined within such a more refined modal space.

The second benefit is not as closely tied to the problem of partiality, but depends on the enhanced logical discrimination. Specifically, this surfaces in the use of a more expressive language, the definability of a wider set of propositions, and finally also on the difference between the factual and the constraining content of a piece of information. In view of the longstanding disagreement on the value of situation-style partiality, it is presumably better to focus on this second kind of theoretical virtue. This is also confirmed by its agreement with the method of abstraction. Even so, the label of situated content which is inspired by Mares' notion of *situated inference* (Mares (2004)) is itself unobjectionable.

DEFINITION 3.2 (SITUATED FACTUAL CONTENT) The situated factual content of A is the set of situations in S (the set of all situations) which do not support A. Hence, we obtain the following generic definition and inequalities which diverge from the classical or worldly account:

$$\begin{array}{lcl} \mathrm{CONT}_s(A) & = & \{s \in S \mid s \nVdash A\} \\ \mathrm{CONT}_s(A) & \neq & \{s \in S \mid s \Vdash\!\sim A\} \\ \mathrm{CONT}_s(\top) & \neq & \emptyset, \text{ for } \top \text{ is a classical tautology.} \\ \mathrm{CONT}_s(\bot) & \neq & S, \text{ for } \bot \text{ is a classical contradiction.} \end{array}$$

This shows among others that on the present situated account the generic notion of exclusion does not longer coincide with the classically equivalent notion of inconsistency. Additionally, since this notion of content is parasitic on the semantics for a relevant logic, contradictions and classical truths are no longer assigned the total and null content. More interestingly, in order to recapture the classical notion of content we do not need to revert to its standard version in terms of maximally consistent models or possible worlds. All we need to refer to, is a special class of situations; the worldly

ones. Indeed, by inspecting the satisfaction-clauses for negation and the functioning of the $*$-operator a restriction can be formulated which cancels the above inequalities. This is done by restricting our attention to situations for which $s = s^*$ holds, for these are the maximally consistent ones or those where $s \nVdash A$ and $s \Vdash \sim A$ express equivalent conditions.[7] Thus, the classical account of content can equally by called the worldly account of factual content, and be defined as:

$$\mathrm{CONT}_w(A) \quad = \quad \{s \in S \mid s = s^* \,\&\, s \nVdash A\}$$

Since this only provides a proper explanation for how we should evaluate the content of formulae which actually do convey factual information (i.e. whose truth-conditions at s are expressed in terms of truth-conditions relative to s itself), we only seem to have provided a proper explication of the content of formulae in which intensional connectives other than negation are not the main connectives. This is where a second notion of situated constraining content becomes relevant.

Roughly, the main consideration is that while $s \Vdash A \oplus B$ or $s \Vdash A \rightarrow B$ do not literally express what is true at s, they surely do convey some valuable content. Such formulae are usually said to confer inferential strength, and with a bit of a hyperbole we could say that the situations which support these formulae function as a constraint that regulates the flow of information between the so connected situations. Presumably, such kinds of content cannot straightforwardly be captured in terms of basic exclusion. By contrast, the idea behind conditional content yields a better starting point. Classically, the content of B that is not yet contained in A is defined as:

$$\begin{aligned} \mathrm{CONT}_w(B \mid A) \quad &= \quad \mathrm{CONT}_w(A \ \wedge \ B) \setminus \mathrm{CONT}_w(A) \\ &= \quad \{s \in S \mid s = s^* \,\&\, s \Vdash A\} \setminus \{s \in S \mid s = s^* \,\&\, s \Vdash A \wedge B\} \end{aligned}$$

An interesting conclusion we can already point out, is that since the worldly content of B relative to A is exactly the worldly content of '$\sim A \vee B$', from a worldly perspective the conditional or relative content of a message is neither purely constraining nor purely factual, but both at the time. This conclusion, however, should not entirely be preserved by a suitable account of situated content, for in that case our intended notion of conditional content should not capture factual content at all. Informally and still too vague for our purposes, the (purely) constraining content of a message $A \rightarrow B$ is given by the proportion of A-cases which are not B-cases. What we need, is a situated precisification of the latter notion; the tools we use to that end are those deployed in Mares (2006) for the definition of a situated probability function. We first define the set of c-accessible A-situations.

$$A^{\sqsubseteq}_c \quad := \quad \{s_i : \exists s(s \Vdash A \,\&\, s \sqsubseteq_c s_i)\}$$

A different way to understand $A^{\sqsubseteq}_c$ is as the set of situations s_i, such that if there is an s which supports A, then $s_i \Vdash A \otimes B$ for all B which are supported by c.

$$A^{\sqsubseteq}_c \quad := \quad \{s_i : \exists s(s \Vdash A \,\&\, \forall B(c \Vdash B \Rightarrow s_i \Vdash A \otimes B))\}$$

This allows us to define the required notion of situated constraining content. In short, it interprets constraining content relative to A as the information effectively added in a c-accessible A-situation, rather than in an A-situation simpliciter.

DEFINITION 3.3 (SITUATED CONSTRAINING CONTENT) The situated constraining content of B relative to A is the content of the c-accessible A-situations which are also B-situations minus the content of the c-accessible A-situations. Namely,

$$\begin{aligned} \mathrm{CONT}_c(B \mid A) &= A^{\sqsubseteq}_c \setminus (A^{\sqsubseteq}_c \cap \{s : s \Vdash B\}) \\ &= A^{\sqsubseteq}_c \setminus \{s : s \Vdash A \otimes B\} \end{aligned} \tag{1}$$

A crucial and enlightening property of this notion concerns its behaviour at logical situations: if $c \in Log$, then if $A \rightarrow B$ is a tautological entailment or relevantly valid implication, it holds that $\mathrm{CONT}_c(B \mid A) = \emptyset$. Namely, $c \in Log$ and $s_1 \sqsubseteq_c s_2$ jointly entail $s_1 \sqsubseteq s_2$, and hence (since situations are provably closed under tautological entailments) every c-accessible A-situation is also a B-situation. Recalling the alleged informational emptiness of logical consequence, we here obtain a characterisation of constraining content that matches the same description. That is, constraints are only informative to the extent that they *exclude* some logically possible extensions of a situation. Obviously, this inverse relationship between informativeness and the logically derivable can only be enforced by taking c to be a logical situation. This should, however, not count as an objection, for that class of situations contains precisely the designated points for the evaluation of theoremhood.

4 INFORMATIVENESS, LOGICAL TRUTH, AND THE TAUTOLOGICAL

The above conclusion that constraining content is tied to relevant entailments in the same sense as worldly content is tied to classical entailments invites us to reconsider the alleged emptiness of logic. While the discussion of this connection was avoided in order to build a pluralist view of content-assessments that did not take the thesis of logical pluralism as a premisse, the topic should not be avoided in general. Two considerations can help us to clarify the scope and meaning of a principle which relates informativeness to what cannot be obtained by purely logico-deductive means.

A first remark derives from the lack of a tight connection between situated factual content and a proper account of logical consequence. In very general terms we might say that the distinction between factual and constraining content largely follows the difference between material and relevant consequence. Of these, only the latter can really be called a logical

kind of consequence (just note that nothing is true at all situations). Our conclusion can therefore only be that situated factual content is hardly logically constrained at all. Even so, this conclusion can not be used as a *reductio* against this kind of content-assessments. Situated factual content is coherent in the sense that it takes the factual content of A to be *null* iff the information that A is available in all situations. It is, in other words, an acceptable but only weakly logically constrained precisification of *GIRP*.

For the second remark we should first take a look at Kratzer's version of situation semantics. The relevant bit of her presentation of a modal space which is more refined than a possible worlds one concerns only the definition of logical consequence and equivalence. She writes: "The notions of 'validity', 'consistency', 'compatibility', 'logical consequence', and 'logical equivalence' depend only on the possible worlds part of propositions. This will ensure that our semantics will be a classical one." (Kratzer (1989): 616). Even if we ignore the actual motives for making this choice, the general point that refinements of modal epistemic space can be conducted independently from the revision of the canons of logical consequences is valuable in its own right. The obvious point it suggests is that for the purpose of doing semantics, the standard level of abstraction suggested by the classical concepts of consequence and equivalence is generally too coarse. More interestingly, it equally implies that a pluralist position about content-assessments, as the one outlined here, need not be accompanied by, or even be dependent on a pluralist position about logical consequence.

This final remark thus confers more weight to our initial choice for not straightly deriving the incompleteness of *GIRP* from the related view Beall & Restall defend about the *Generalised Tarski Thesis* combined with additional assumptions about the relation between logical consequence and informativeness. It also motivates the digression about the virtues of the method of abstraction, and the subsequently defended generality of content-assessments within a modelling-practice. By taking this road we have not only argued for the LoA-dependence of content-assessments, but also formulated a pluralist thesis about informational content which does not necessarily and entirely denies logical orthodoxy.

NOTES

1. Put simply, the basic cases are those where the update or the information an agent receives is of a factual nature. This excludes an input containing knowledge-claims, and more generally any kind of message containing so-called epistemic modals. Formally, the class of factual messages can be reduced to those expressible with negation and conjunction.

2. For instance, if we say that the content of p is the set of *cases* it is inconsistent with, it does not make sense to interpret the notion of a *case* as a partial and possibly inconsistent situation.
3. A view I previously defended by endorsing the thesis that "the information that A is only informative to the extent that A lies outside the scope of logical consequence alone."
4. There are of course exceptions to this rule, for instance when one introduces *semantic blindness* to explain why an otherwise competent speaker fails to abide by the formal semantics one takes to be right. Yet, the fact that the need to defer to semantic blindness is often considered a theoretical flaw, clearly enforces the above point.
5. The standard notation as Rcs_1s_2 is (following Fine (1974)) rewritten as $s_1 \sqsubseteq_c s_2$ to make the connection with $\sqsubseteq$ more explicit. A comparable notation is Restall's channel-based version: $s_1 \overset{c}{\mapsto} s_2$ (Restall (1995)).
6. $\wedge$ and $\vee$ are reserved for the classical connectives.
7. Likewise, the combination of $\sqsubseteq$ and $*$ can also be used to delimit the set of consistent, but possibly partial situations, viz. $\{s \in S \mid s \sqsubseteq s^*\}$.

REFERENCES

Baltag, A. and L. S. Moss, 2004, Logics for Epistemic Programs, *Synthese*, 139(2): 165–224.

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

Barwise, J. and J. Perry, 1999, *Situation and Attitudes.* CSLI Publications, Stanford.

Beall, J. C. and G. Restall, 2006, *Logical Pluralism.* Oxford University Press, Oxford.

Carnap, R. and Y. Bar-Hillel, 1952, An Outline of a Theory of Semantic Information, MIT, Cambridge, Massachusetts, Technical Report 247. Reprinted in Bar-Hillel, Y., 1964, *Language and Information. Selected Essays on Their Theory and Application.* Addison-Wesley, Reading / London.

Dunn, J. M., 1976, Intuitive Semantics for First-Degree Entailments and 'Coupled Trees', *Philosophical Studies*, 29:149–168.

Fine, K., 1974, Models for entailment, *Journal of Philosophical Logic*, 3(4): 347–372.

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

Floridi, L., 2006, The Logic of 'Being Informed', *Logique & Analyse*, 49(196): 433–460.

Floridi, L. and J. W. Sanders, 2004, Levellism and the Method of Abstraction, Information Ethics Group - Oxford University.

Humberstone, I. L., 2005, Logical Discrimination, in: *Logica Universalis*, J.-Y. Béziau, ed., Birkhäuser Verlag, Basel: 207–228.

Israel, D. and J. Perry, 1990, What is Information, in: *Information, Language and Cognition*, P. Hanson, ed., University of British Columbia Press, Vancouver: 1–19.

Kemeny, J. G., 1953, A Logical Measure Function, *The Journal of Symbolic Logic*, 18(4): 289–308.

Kratzer, A., 1989, An investigation of the lumps of thought, *Linguistics and Philosophy*, 12(5): 607–653.

Mares, E., 1997, Relevant Logic and The Theory of Information, *Synthese*, 109(3): 345–360.

Mares, E., 2004, *Relevant Logic. A Philosophical Interpretation.* Cambridge University Press, Cambridge.

Mares, E., 2006, Relevant Logic, Probabilistic Information, and Conditionals, *Logique & Analyse*, 49(196): 399–411.

Priest, G., 2001, Logic: One or Many?, in: *Logical Consequence: Rival Approaches*, J. Woods and B. Brown, ed., Hermes, Stanmore: 23–38.

Read, S., 2006, Monism: The One True Logic, in: *A Logical Approach to Philosophy. Essays in Honour of Graham Solomon*, D. De Vidi and T. Kenyon, ed., Springer, Dordrecht: 193–209.

Restall, G., 1995, Information Flow and Relevant Logic, in: *Logic, Language and Computation: The 1994 Moraga Proceedings*, J. Seligman and D. Westerståhl, ed., CSLI Press, 463–477.

Sequoiah-Grayson, S., 2006, Information Flow and Impossible Situations, *Logique & Analyse*, 49(196): 371–398.

Shannon, C. E. and W. Weaver, 1998, *The Mathematical Theory of Communication.* University of Illinois Press, Urbana / Chicago.

Van Benthem, J., 2006+, Information as Correlation vs. Information as Range, in: *Memorial Volume for Jon Barwise*, L. S. Moss, ed.