

# Hard and Soft Logical Information

## Abstract

In this paper I use the distinction between hard and soft information from the dynamic epistemic logic tradition to extend prior work on informational conceptions of logic to include non-monotonic consequence-relations. In particular, I defend the claim that at least some non-monotonic logics can be understood on the basis of soft or “belief-like” logical information, and thereby question the orthodox view that all logical information is hard, “knowledge-like”, information.

## 1 INTRODUCTION

Logic deals with specific classes of truths and truth-preserving relations, namely the necessary and formal ones. Within an informational conception of logic, this claim translates to the thesis that all logical information is hard information—the kind of information we more readily associate with knowledge than with belief. The upshot of this paper is to develop an account of logic within which we can theorise about hard, but also about soft logical information—the belief-like logical information that shapes ampliative inference. The resulting inflation of what counts as logical information is not meant to deny the distinction between deductive or truth-preserving and ampliative or merely plausible styles of inference, but instead proposes a philosophical account of logical consequence that accommodates several styles of inference (or indeed, informational practices) within a single conceptual framework. This goal is worthwhile on its own (and fills a void in the current literature),<sup>1</sup> but it is also indispensable if we want the philosophy of logic to remain relevant to the wide variety of practices in contemporary logic.

---

<sup>1</sup>Existing work in this domain generally addresses different questions, like the study of the structural properties of nonmonotonic consequence (Gabbay 1994), the place of defeasible inference in our epistemic and practical lives (Gabbay and Woods 2003; 2008, Horty 2012), or the formal integration of the diverse systems of non-monotonic logic (Makinson 2005, Strasser 2014).

As a contribution to the philosophy of logic and the philosophy of information, the present paper does not as such contribute novel formal results about existing logical systems nor does it present any new systems. Instead, it uses known results to develop a philosophical account of logical consequence, and in particular does so by relying on insights from logics for multi-agent systems, dynamic epistemic logic, and non-monotonic logics that are only very rarely taken into account in philosophical discourses and are (unfortunately) largely absent from debates on the nature of logical consequence.

Before the proposed account of hard and soft logical information can even be formulated, a substantial amount of groundwork must be laid, and the broader context of the present proposal must be sketched.<sup>2</sup> In particular, I should clarify what I mean by an informational conception of logic, and explain how I'll use the concepts of hard, soft and logical information, and fit these into the broader framework. This material is covered in Sections 2 (the basics), 3 (hard and soft information), and 4 (logical information). These three sections are largely independent of each other, and come in support of Section 5 where I defend the central thesis that at least some soft information can be described as a kind of logical information because it encodes norms or criteria for consistency, non-triviality and consequence, and provides the basis for reliable informational interaction.

## 2 AN INFORMATIONAL ACCOUNT OF LOGIC

Information is a notoriously polymorphic concept, and explaining how it relates to logic is an equally controversial issue. As used here, information is best understood in connection to (and as the outcome of) successful modelling, where a model is the result of adopting a level of abstraction that lets us focus on an explicitly defined set of features, called *observables*, of the system under consideration. If we want to explain the relation between logic and information on the basis of this idea, it is convenient to see

- modelling as the process of situating a system within a logical space, and
- a logical space as an abstract structure that formalises the distinctions afforded by the adoption of a level of abstraction.

Truth-tables for propositional logic are one of the most basic examples of such a logical space. If the *system* we want to model is some (possibly non-strict) part of reality and we care only about whether in this part of reality  $p$  and  $q$  hold, then

---

<sup>2</sup>Given the substantial deviations from current logical orthodoxy, I have had to sacrifice brevity and conciseness to eliminate external dependencies.

the relevant logical space can be represented as a 4-line truth-table. Situating the system within that logical space can then be understood as the selection of one or more lines of that truth-table that are in agreement with the data (understood as *constraining affordances*; see Floridi 2011*b*, 76) obtained from the system.

From a formal point of view, it is also useful to extend this proposal with an explicit distinction between the meta-language we use to characterise a logical space (often a set-theoretic language) and the object-language we use to formulate theories about a given system or to describe individual models of that system. In some cases, this implies that not every subset of a logical space can be exactly described by the object-language, or that two distinct elements of a logical space are indistinguishable from the standpoint of the object-language.

From a philosophical point of view, it should then be emphasised that when we adopt a level of abstraction, and situate a system within the logical space afforded by that level of abstraction, we are in the business of *attributing* properties to a system (Floridi 2011*b*, Chapters 3, 8 and Floridi 2011*a*). As a consequence, we *posit* possible states of that system as well as the corresponding distinctions between these states in accordance with the level of abstraction we adopted. We may therefore say that a logical space supplies the propositions or properties in terms of which a system can be analysed, whereas the object-language determines which of these properties can effectively be expressed and thus used to formulate a theory or reason about a system. The constraints imposed by the language on what can be expressed are, in this perspective, pragmatic (which properties/propositions can be used in our theorising) rather than semantic or even ontological (which properties/propositions exist in the logical space).

Given a logical space and an object-language, it is easy to see how we can think of a logical consequence-relation defined over this object-language (and interpreted in that logical space) as the logic that is associated with a given modelling practice. That is, we can think of that logic as a set of standards or criteria for consistency, non-triviality (two features of successful modelling), and valid deduction. Figure 1 (a generalisation of Floridi 2008, 230) below summarises the basic concepts used above to connect logic to information.

The resulting perspective on logic agrees with how logical tools and methods are often used, but departs from existing “logic-as-modelling” views (see Cook 2010, §6 and Shapiro 2014, Chapter 2) in the sense that it does not take the modelling of certain human linguistic practices, like deductive and argumentative practices that use natural language, as the primary goal of any logical theorising. In particular, the current proposal departs from the traditional view that logic should always be

## Hard and Soft Logical Information

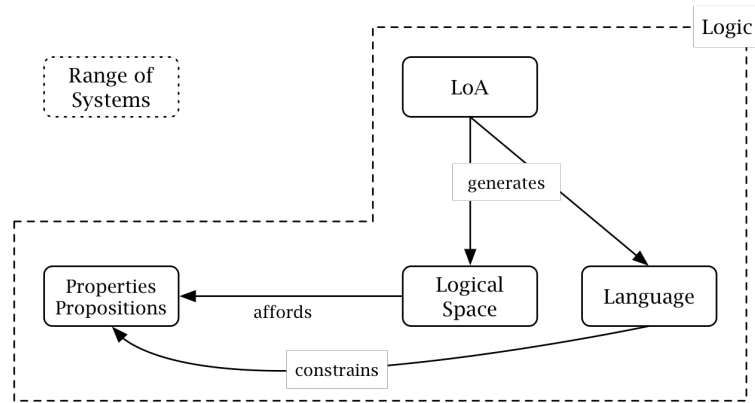


FIGURE 1: Levels of abstraction and logical spaces

understood and evaluated relative to such existing natural-language practices.<sup>3</sup> It does not, however, imply that natural language practices cannot successfully be modelled, or that such modelling cannot be philosophically or scientifically fruitful.

When we think of the adoption of a level of abstraction as the specification of a set of relevant observables, understood as predicates with their intended interpretation, it seems obvious to identify levels of abstraction with the extra-logical part of a language (the *signature* of a language) and the intended interpretation of this part of the language. This is also the view presupposed by the earlier truth-table example, and suffices to capture the fact that a language with more predicates is

1. more expressive (more can be said about a system), and
2. more discerning (more possible states of the system can be distinguished).

We can refer to (1) and (2) as *informational virtues*, features of a level of abstraction that give us a better grip on the system we try to model.



Informational virtues are similar to theoretical virtues in the sense that, all else being equal, more expressive and more discerning models and theories are preferable. Other theoretical virtues like simplicity and parsimony, or indeed the need for tractable logical systems, may, however, exert an opposite pull, and hence favour the inclusion of fewer observables or the adoption of a less expressive logical system

<sup>3</sup>This suggestion goes against a long tradition in the philosophy of logic, but can be supported by recent work of Michael Glanzberg that underscores the substantial abstraction and idealization needed to connect the semantics of natural language to the model-theory that supports logical consequence relations (Glanzberg 2015).

(see also below). This gives rise to trade-offs between informational, computational and other theoretical virtues, and it is only by taking into account the purpose of a given inquiry that we can decide which virtues to favour and hence which level of abstraction to adopt.

Expressiveness and discernibility do not only vary with the available observables (or, equivalently, with the signature of a language), but also with the purely logical part of the language we use to describe models. Several examples of this fact, which makes the informational perspective on logic thoroughly pluralist, can be given. A textbook example is how first-order quantification drastically extends the expressiveness of zero-order languages. More interestingly, sub-classical logics also introduce distinctions that are beyond the reach of classical logic: Paraconsistent logics keep apart the different inconsistent theories that classical logic identifies with the only inconsistent theory it can accommodate, namely the trivial theory, whereas intuitionistic logic affords fine distinctions between constructively true and merely weakly true theorems (i.e. doubly negated theorems) of classical mathematics.<sup>4</sup> At the more abstract level of logical spaces it is clear that with more truth-values (one of many departures from classical logic) we immediately obtain more possible assignments (think again about the lines in a truth-table) and hence the possibility to entertain more possible states of the system under consideration.

Looking at distinctions afforded by a sub-classical logic by examining how it shapes a logical space is doubly interesting. First, because it makes it explicit that what counts as a valuable distinction depends on the system we're looking at as well as on our reasons for investigating that system. On this account, we adopt a given logical space (and a language) because it provides us the conceptual (and semantic) means to reason about the possible states of the system we have no good reason to rule out. A second way of looking at the distinctions and possible states afforded by a logical space circumvents the worry that our knowledge of the possible states of a system is itself mediated by the level of abstraction we adopt, and points to a crucial trade-off between the ability to make distinctions and the ability to draw conclusions. Because richer logical spaces make room for finer distinctions, they inevitably also make room for additional counterexamples. This means that finer distinctions always come at a price. We have to give up certain inference forms as well. With this in mind, it shouldn't surprise us that the non-equivalence of  $p$  and  $\neg\neg p$  in intuitionistic logic, or the non-equivalence of  $p \wedge \neg p$  and  $q \wedge \neg q$  in many

---

<sup>4</sup>Shapiro (2014, Chapter 3) describes the more elaborate example of Kock-Lawvere's smooth infinitesimal analysis, a mathematical theory that is classically inconsistent (and hence trivial), but intuitionistically consistent. (This is an interesting example because since *Ex Falso* is intuitionistically valid, the relevant distinction is not a matter of paraconsistency, but of preventing the inconsistency to arise in the first place.)

paraconsistent logics are literally interchangeable with the invalidity of, respectively, *reductio ad absurdum* and *ex contradictione quodlibet*.

These facts about logical spaces establish an inverse relationship between the deductive strength and discriminatory power of consequence relations defined over logical spaces, which in its turn suggests that the adoption of a logic (and hence the adoption of a level of abstraction, or the construction of a logical space—the interchangeability of these three should be evident by now) is primarily based on the requirement to coordinate our need to draw fine distinctions, and our need to reach interesting conclusions. Since this is more a pragmatic (i.e. purpose-relative) than an ontological matter, the question whether a point in logical space matches a possible state of the system loses much of its appeal. Even core logical requirements like consistency are reduced to their bare minimum, namely absolute consistency or non-triviality, which can either be motivated as a standard for successful construction or as a corollary of the requirement that not all distinctions must be sacrificed to reach more conclusions.

Logical spaces carve out informational contents, but these cannot at the same time be *contentfull* and *discerning*. As a consequence, we cannot adopt a logic without having to negotiate the trade-off between these two intuitively desirable, but formally opposed informational virtues. Except for the two extremes (the trivial and the empty logic, which correspond to an empty logical space and the set of all *open* worlds), any option may be the right one for a given purpose. Although in the abstract neither informational virtue is more important than the other, the best level of abstraction need not be somewhere in the middle. This is why it is better to say these virtues need to be negotiated rather than merely balanced, as the latter implies a symmetry that is absent in most concrete situations. The more we know about the system (the constraining affordances) and about our reasons for investigating that system (the purpose), the easier it should become to break this symmetry.<sup>5</sup>



The metaphor of logical space is mathematically and conceptually very powerful, but it is also philosophically suggestive (Stalnaker 1984, Rayo 2012). Indeed, talk about logical spaces without a good story about its metaphysics is often considered highly problematic. When it comes to the development of an informational semantics, logical spaces are best seen as an abstraction of how we access and use information (Allo and Mares 2012); a mathematical model of how available distinctions lead to

---

<sup>5</sup>From a pluralist perspective logic choice need not even be definitive or exclusive; not only can choices be revised, but multiple logics could indeed be used in parallel for certain types of inquiry. The description of the roles of first and second-order logic in the foundations of mathematics, as described by Shapiro (2000), is an example of the latter case.

a consequence relation that can fulfil its theoretical role without further metaphysical support.<sup>6</sup> The availability of a complementary purely syntactic perspective on the inverse relation between deductive strength and discriminatory power provides further evidence for the view that logical spaces need a logical or mathematical, but not a metaphysical foundation.

The required connection can be described with the help of synonymy-relations. This is a generalisation of equivalence-relations that was originally devised in Smiley (1962) as a criterion for the definability of a connective in a given language. More recently, Humberstone (2005) repurposed it to reason about logical discrimination. Synonymy, in other words, is a criterion for sameness of information (in virtue of logic) of formulas.<sup>7</sup> Thus, we say that the expressions  $A$  and  $B$  are synonymous relative to  $\vdash$  (henceforth,  $A \equiv_{\vdash} B$ ) iff:

$$C_1(A), \dots, C_n(A) \vdash C_{n+1}(A) \text{ iff } C_1(B), \dots, C_n(B) \vdash C_{n+1}(B), \quad (\text{Syn}_{\vdash})$$

where each  $C_i(B)$  is the result of replacing zero, one, or more occurrences of  $A$  in  $C_i(A)$  by  $B$ . On this account, synonymous expressions are expressions that can freely stand in for each other without thereby affecting what can and cannot be derived. As shown by Humberstone, the formal connection between deductively stronger logics and the fudging of distinctions

$$\vdash_0 \subseteq \vdash_1 \text{ iff } \equiv_{\vdash_0} \subseteq \equiv_{\vdash_1}$$

holds for a wide range of logics.<sup>8</sup>

On the account of logic presented in this section, a logic is an inherent part of a level of abstraction because it co-determines which distinctions can be made. Conversely, the adoption of a level of abstraction cannot be dissociated from how we choose a logic because available distinctions impact what follows from what. This co-dependence is of particular importance when logical discrimination and deductive strength are inversely proportionate, since in those cases the choice of a level of abstraction must be the result of negotiating between two important informational virtues, namely the contentfulness and discerning nature of information.

<sup>6</sup>Compare with “I take abstract entities to be human constructions, the ability to create such constructions being made possible by symbolic artefacts such as language and mathematics.” (Giere 2004, fn. 7).

<sup>7</sup>This generalises one of the requirements of adequacy for a theory of semantic content in Carnap and Bar-Hillel (1952, 7), namely that all L-equivalent sentences should be assigned the same information.

<sup>8</sup>The typical case is one where synonymy and logical equivalence are co-extensive, but counterexamples can be found in substructural logics with equivalence-relations that are coarser than their synonymy relations.

3 HARD AND SOFT INFORMATION

Characterisations of logical consequence traditionally rely on a modal vocabulary—even if its extension can be characterised in purely extensional terms—and see the modal profile of logical truth and logical consequence as one of its most distinctive features: logical truths are necessary truths, and logical consequence requires necessary truth-preservation. Using a more careful formulation, we could say that logical consequence (and truth) deals with relative (and absolute) strict necessities that hold in virtue of their logical form. Spelling out the nature of the underlying modality (alethic, logical, metaphysical) as well as the precise way in which we should understand the formality-requirement is particularly hard and the current debate leaves much room for controversy.<sup>9</sup> By comparison, the association with strict modalities is, exactly like the association of logic with deductive inference I mentioned earlier, less controversial; especially when necessity is understood formally in purely quantificational terms rather than substantially.<sup>10</sup> Indeed: necessity is just what makes the conclusions of deductive arguments unavoidable, and hence also what separates deductive from non-deductive inference.

Because modal distinctions are hard to use in an entirely neutral manner, I will here adopt the concepts of *hard* and *soft* information as a means to clarify the distinction between the deductive and non-deductive realms of logic. This distinction is commonly used in the field of dynamic epistemic and doxastic logics, where it is used as a high-level distinction between knowledge-like and belief-like information (Baltag and Smets 2008, van Benthem 2007). Here, it is attractive because of its emphasis on information, but also because it unifies several features that separate the deductive from the non-deductive. In particular, it captures the distinctions between strict and variable quantification and necessity, between monotonicity and non-monotonicity,<sup>11</sup> and between static and dynamic aspects of information. Hard information is information which, when held, can only be expanded, and, when received, can only expand our information.<sup>12</sup> Analogously, soft information is infor-

<sup>9</sup>Recent discussions of the modal aspect can be found in Rumfitt (2015, Chapter 3) and Cresswell et al. (2016). On formality, see MacFarlane (2000), Dutilh Novaes (2011).

<sup>10</sup>On the distinction between substantive and non-substantive modality, see Rini (2013). The emphasis on necessity in Beall and Restall (2006) is in my view of the more formal kind, whereas the objections against Tarski's characterisation of logical consequence in Etchemendy (1990) requires a more substantial view on necessity.

<sup>11</sup>In many cases, truth-preservation and monotonicity come and go together, but as shown by Restall (2016) some nonmonotonic consequence-relations are truth-preserving. Applied to the present context: some soft information—like *safe belief*—is more knowledge-like.

<sup>12</sup>It is tempting to see Moorean phenomena, the so-called paradox of serious possibility, and the existence of unsuccessful updates (which can be linked to Fitch's paradox) as counterexamples to



mation which, when held, can still be revised in view of new information, and, when received, is (often) conservatively aggregated with rather than dictatorially imposed on one's information.

Strict and variable modalities, and thus also hard and soft information, can be distinguished at multiple levels. They can be defined at the meta-level by taking a set of models (of a language) to be the relevant logical space. In that case we identify the strict modalities with logical truth and logical consequence (truth, resp. truth-preservation, in all models), and a relative variable modality with truth-preservation in a selection (based on an ordering) of all models. They can also be defined at the object-level by taking a single model (of a modal language) to be the relevant logical space, with Kripke-style models as the paradigmatic example. In that case, necessity and strict implication capture hard information, and various conditionals capture soft information. Non-monotonic consequence (as opposed to various non-monotonic conditionals) is a meta-theoretical notion, but can also be characterised in the object-language. Technically, this has the advantage that we can characterise all the relevant modalities in Kripke-style structures, while philosophically it has the advantage that we can, first, rely on the analogy between knowledge and belief, and hard and soft information, and, second, that we can distinguish between logical and extra-logical hard and soft information.



To develop this further, let us take a combined modal logic of knowledge and belief as a starting point. If we assume that knowledge and belief are both positively and negatively introspective, and only differ in the fact that we can have false beliefs but no false knowledge, knowledge can be characterised as an **S5** modality and belief as a **KD45** modality. When complemented with principles that relate knowledge with belief (e.g. to ensure that we believe everything we know), we obtain a combined logic for highly idealised notions of knowledge and belief (Gochet and Gribomont 2006, §2). Since our present focus is the principled distinction between hard and soft information, rather than the development of corresponding cognitively realistic or sufficiently fine-grained accounts of information, several well-known problems of this formalisation can provisionally be ignored. What I care about is an inherent limitation of the characterisation of belief with **KD45** operators. This is best explained by introducing a first basic type of Kripke-models.

**DEFINITION 1 (KB STRUCTURE)** A KB structure is a 4-tuple  $\mathfrak{M} = (S, \approx, \rightarrow, \|\cdot\|_{\mathfrak{M}})$ , with these features. This is mistaken, as the latter are essentially problems for purely syntactic characterisations of the information held by an agent (especially the belief-sets presupposed by AGM-style belief-revision theory).

## Hard and Soft Logical Information

$S$  a set of points,  $\|\cdot\|_{\mathfrak{M}}$  a valuation-function, and  $\approx$  and  $\rightarrow$  binary relations over  $S$  such that (i)  $\approx$  is an equivalence relation, (ii)  $\rightarrow$  is serial, transitive and euclidean, and (iii)  $\rightarrow \subseteq \approx$ .

The following standard satisfaction-clauses are then sufficient to ensure that knowledge ( $[\approx]$ ) and belief ( $[\rightarrow]$ ) have exactly the features we put forward:

$$\mathfrak{M}, s \Vdash [\approx]A \text{ iff } [s]^\approx \subseteq \|A\|_{\mathfrak{M}} \quad (1)$$

$$\mathfrak{M}, s \Vdash [\rightarrow]A \text{ iff } [s]^\rightarrow \subseteq \|A\|_{\mathfrak{M}} \quad (2)$$

with  $[s]^\approx = \{t \in S \mid s \approx t\}$  and  $[s]^\rightarrow = \{t \in S \mid s \rightarrow t\}$ , namely the semantic propositions that correspond to an agent's knowledge and belief. Remark finally that the restriction of  $\rightarrow$  to  $[s]^\rightarrow$  is an equivalence-relation even though  $\rightarrow$  isn't in general reflexive.

On this account, knowledge and belief are both captured by a strict modality: the clauses (1) and (2) are both based on a universal quantification over accessible states. If we want to characterise actual knowledge and belief, this is all we need: the relations  $\approx$  and  $\rightarrow$  encode exactly the information we need for this purpose. When it comes to belief, this picture is however unduly restrictive, as  $\rightarrow$  does not encode the revision-policy that is needed to account for the revisability of belief. Like  $\approx$ , the  $\rightarrow$  relation only encodes hard information. This clashes with the suggestion that belief is a kind of soft information.

The above diagnosis leads us to a core feature of soft information: soft information extends our hard information by further refining it. On the traditional account of knowledge and belief, however,  $[s]^\rightarrow$  restricts  $[s]^\approx$  (and thus carries more information) without refining it. A more general type of Kripke-model (based on Baltag and Smets 2008) allows us to define  $[s]^\rightarrow$  as a refinement of  $[s]^\approx$ .

**DEFINITION 2 (HARD/SOFT STRUCTURE)** An information structure for hard and soft information is an 4-tuple  $\mathfrak{M} = (S, \approx, \leq, \|\cdot\|_{\mathfrak{M}})$ , with  $S$  a set of points,  $\|\cdot\|_{\mathfrak{M}}$  a valuation-function, and  $\approx$  and  $\leq$  binary relations over  $S$  such that  $\approx$  is an equivalence relation and  $\leq$  is a pre-order that refines  $\approx$  ( $s \leq t$  only if  $s \approx t$ ).

The hard information available at  $s$  is defined as before by  $[s]^\approx$ , but the relation  $\rightarrow$  we used to characterise belief can now be defined in terms of the pre-order  $\leq$ :

$$s \rightarrow t := t \in [s]^\approx \ \& \ \forall t' \in [t]^\approx \ (t \sim t' \implies t \leq t')$$

with  $\sim := \leq \cup \geq$ . That is,  $s \rightarrow t$  iff  $t$  is  $\leq$ -minimal in  $[s]^\approx$ . If we furthermore assume that  $\leq$  is connected and well-founded within any  $\approx$ -cell (both standard assumptions of plausibility-models),  $\rightarrow$  is indeed serial, transitive and euclidean as in

a KB-structure. Moreover, given these assumptions,  $\approx$  can itself be defined as the symmetric closure of  $\leq$ , and thus  $\leq$  can be seen as an encoding of both hard and soft information (which confirms my point that soft information is best seen as a refinement of hard information).

The strict relative modality (i.e. the strict implication of C.I. Lewis) is definable in the usual way

$$A \multimap B := [\approx](A \supset B),$$

but  $[\rightarrow](A \supset B)$  is not a genuine variable relative modality. Instead we need:

$$\mathfrak{M}, s \Vdash [\xrightarrow{A}]B \text{ iff } [s] \xrightarrow{A} \subseteq \|B\|_{\mathfrak{M}} \quad (3)$$

with

$$s \xrightarrow{A} t := t \in [s] \approx \cap \|A\|_{\mathfrak{M}} \ \& \ \forall t' \in [t] \approx \cap \|A\|_{\mathfrak{M}} \ (t \sim t' \implies t \leq t')$$

and  $[s] \xrightarrow{A}$  defined in the same manner as before.

This fact is supported by multiple impossibility-results (see e.g. Leitgeb 2007), but the basic idea can already be appreciated by considering that  $[\rightarrow](A \supset B)$  a claim about what is actually believed that cannot be used to pre-encode dynamic changes. Indeed, whereas  $[\approx](A \supset B)$  is itself also a claim about what is actually known, it can also be used to pre-encode what will be known after one learns  $A$ .<sup>13</sup>

To summarise: (i) hard information is characterised semantically in terms of an unstructured set of options encoded by an equivalence-relation, can be captured syntactically by strict modal operators, and has a monotonic dynamics; and (ii) soft information is characterised semantically by a structured set of options encoded by a pre-order, can be captured syntactically by a range of variable modal operators, and has a non-monotonic dynamics. While it can be useful to think of hard information as knowledge-like information, and soft information as belief-like information, the main moral of this section is that one should think of soft information as a refinement, and not just as an extension of one's hard information.

#### 4 LOGICAL INFORMATION

In a popular use of the term, “logical information” refers to the information we obtain through logical inference. Spelling out what kind of information it is typically requires one to develop a suitable notion of fine-grained content (Jago 2012) or otherwise solve the scandal of deduction (D’Agostino and Floridi 2009, Rayo 2012).

<sup>13</sup>The formal counterpart of this claim is that reduction-axioms for changes in hard information do not (except for the case of common knowledge, see van Benthem 2011, §3.5) require a dedicated conditional construction, but changes in soft information do.

As I'll use the term, it refers to the information we have in virtue of logic; a use that is much closer to the concept of a logical situation in the semantics of relevant logic (see especially Mares 2010, §9) or normal worlds in the semantics of non-normal modal logics. It is also a use that doesn't require one to deal with the scandal of deduction: we can talk about logical information and about the logical information we have or have access to, but also accept that the information yield of logical deduction is null.

Formally, and especially if we develop our theory with the help of the metaphor of logical space, we have a rather good idea of what logical information is. It is the information we associate with the total logical space.<sup>14</sup> On a standard account of informational content this is equivalent to saying that it is the empty information, or the information associated with the empty premise-set.

Despite the overall pessimism that surrounds the scandal of deduction, thinking about the nature of logical information can help us to further clarify how logic and information are related. As a provisional definition, we can stipulate that logical information is the information we have solely in virtue of the level of abstraction we adopt. That is, logical information is information afforded by the available distinctions. Depending on one's tendency, this description may seem risky (conceptual knowledge in disguise?) or just evasive because it does not really offer an explanation that exceeds the formal framework (be it the metaphor of logical space, or just synonymy-relations). The pluralist and constructionist basis of the informational conception of logic from Section 2 is nevertheless sufficient to dismiss the first worry: concepts (or distinctions) are not 'given' but developed as part of our modelling activity.<sup>15</sup> If our modelling activity is successful, then its underlying logic will be correct or adequate for the purpose at hand.<sup>16</sup>

---

<sup>14</sup>This description does not straightforwardly apply to logical spaces that include both logical and non-logical points. The guiding idea can nevertheless be maintained since even when logical information is not uniformly available in the logical space, we still look at the total logical space from the perspective of its logical fragment (formally: consequence and theoremhood are evaluated at a logical point).

<sup>15</sup>The following remark taken from recent work of Jeremy Avigad is particularly illustrative because it uses a similar vocabulary: "Note that I am using the word "language" here in a very broad sense to refer not only to our choice of words and grammar, but also the way we frame, conceptualize, analyse, and reason about the problems before us. What I am advocating is a view of mathematics as a linguistic artefact, something we have designed, and continue to design, to help us get by in the world. (...) Various pieces of mathematics are designed to help us achieve various goals, and like any human artefacts, they can serve their purpose well or poorly. The philosophical challenge is to understand the general principles by which they serve our goals well, and to support the development of better artefacts. In that sense, the philosophy of mathematics can be viewed as a kind of language engineering, akin to any other design science." (Avigad 2015, §1).

<sup>16</sup>The reader familiar with the distinction between *strong* and *weak* endorsements of a logic (Beall

Our provisional definition of logical information can be made more robust by thinking of logical information as the most basic or general information that shapes a given modelling activity by supplying criteria for consequence, consistency and non-triviality, *and* by providing the basis or even the condition of possibility for reliable interaction. In other words, a common level of abstraction ensures reliable communication about the features of a given system. This can be illustrated by the fact that communication across levels of abstraction is often (and especially when it isn't simply cross-purpose) equivocal or even noisy.<sup>17</sup>



A general account of this proposal is more effectively formulated by analogy with the notions of common knowledge and common information: logical information is common information in virtue of the distinctions afforded by a level of abstraction, and this not only guarantees reliable interaction but also grants certainty about reliable interaction. Making this point requires one to adopt a multi-agent and interactive perspective on logic that goes against the tradition of (most of) 20<sup>th</sup> century logic, and sees the single-agent case as the limit of the multi-agent case instead of seeing the multi-agent case as an extension of the single-agent case.<sup>18</sup>

Common hard and soft information are easy to define in a multi-agent extension of the Kripke-structures we described in the previous section. Let us, provisionally, focus on common hard information (or common knowledge).

**DEFINITION 3 (MULTI-K STRUCTURE)**

A multi-K structure is a 4-tuple  $\mathfrak{K} = (S, A, \{\approx_i\}_{i \in A}, \|\cdot\|_{\mathfrak{K}})$  with  $S$  a set of points,  $A$  a set of agents, each  $\approx_i$  an equivalence relation over  $S$ , and  $\|\cdot\|_{\mathfrak{K}}$  a valuation function.

Given such a structure, we can define the relation  $\approx_A$  as the symmetric and transitive closure of  $\bigcup_{i \in A} \approx_i$ , which is equivalent to saying that  $s \approx_A t$  iff there are finite

and Restall 2006) can now fully appreciate what we won by replacing talk about actual states of the system (or genuine possibilities) with the need to negotiate between informational virtues. I leave it open whether this approach can fully resolve disputes over the pragmatics and truth of how a logical space is constructed (Hofweber 2014, Rayo 2014; 2012).

<sup>17</sup>One example of this, developed in Allo (2013c), is that from the perspective of many non-classical logics (and assuming a similar language) classical logic is a noisy logic: missing distinctions do not surface as uncertainty over the message that is sent, but as noise. If, for instance, we only have the distinctions afforded by classical logic, we can never convey the intuitionistically weaker  $\neg\neg p$  without also conveying the stronger  $p$ .

<sup>18</sup>The multi-agent perspective is richer, and allows one to highlight features of logic that disappear or become vacuous when only a single agent is taken into account (this too is the method of abstraction at work!). In Allo (2013b) I used a similar strategy to think about closure and introspection. Another example is Dutilh Novaes's (2015) adoption of an explicit dialogical view on logic to discern a normative aspect of logic that is lost when one focuses on the reasoning of an individual.

sequences of points  $s_1, \dots, s_n$  and agents  $a_1, \dots, a_{n+1}$  such that  $s \approx_1 s_1$ ,  $s_m \approx_{m+1} s_{m+1}$  for each  $m < n$ , and  $s_n \approx_{n+1} t$ .

Since for any binary relation  $\approx_i$  we can define a unary operator  $[\approx_i]$ , we can define such an operator for  $\approx_A$  as well and this captures the established notion of common hard information. The formal properties of this type of operator are well understood (it's a fixed point, and is not equivalent to any finite iteration of "everybody is informed"). There is, however, one formal result that can be used to explain why common hard information and information about the total logical space are perfectly interchangeable from the perspective of all the agents in  $A$ . Let  $\mathfrak{M}_1 = (S_1, A, \{\approx_i\}_A, \|\cdot\|_{\mathfrak{M}_1})$  and  $\mathfrak{M}_2 = (S_2, A, \{\approx'_i\}_A, \|\cdot\|_{\mathfrak{M}_2})$  be such that:

- $S_2 = \{s \in S_1 \mid s_0 \approx_A s\}$ ,
- $\{\approx'_i\}_A = \{\approx_i\}_A \cap S_2 \times S_2$ , and
- $\|\cdot\|_{\mathfrak{M}_2}$  is the restriction of  $\|\cdot\|_{\mathfrak{M}_1}$  to  $S_2$ .

Hence,  $\mathfrak{M}_2$  is just the result of reducing  $\mathfrak{M}_1$  to what is commonly known at  $s_0$  in  $\mathfrak{M}_1$ . As a generated submodel of  $\mathfrak{M}_1$ ,  $\mathfrak{M}_2$  is bisimilar and hence also modally equivalent to  $\mathfrak{M}_1$  (Blackburn et al. 2001, §2.1, 2.2). This implies that as far as the agents in  $A$  are concerned, there is no difference at all between the two models:<sup>19</sup> there is no information about the logical space or room for distinctions over and above what is afforded by the common hard information encoded in the logical space.

The above argument explains why nothing is lost by replacing an unrestricted total logical space by the common hard information in that space, but does not yet establish that the *common ground* role of logical information leads to certainty of reliable interaction. In the context of a dialogue, and assuming an intuitionist background logic, if Alice puts forward  $\neg\neg p$  she knows that she will not have to defend the stronger claim  $p$ , and also knows that if her interlocutor Bob later claims that  $p$  he will not solely do so in virtue of her prior claim. That is, Alice knows that she can accept Bob's claim as new information, but Bob also knows that she will accept his claim as novel information. In general, they are both confident (and this fact is common information as well) that throughout their exchange no information will be created *ex nihilo* or merely due to noise caused by misaligned logical distinctions. In an  $n$ -agent interaction this is even more beneficial, for there it will ensure that even if Alice and Bob have perfectly aligned logical distinctions, their reliance on Carol as an intermediary to pass on  $\neg\neg p$  will not introduce  $p$  in the communication unless this new information is independently made available to Carol. Further generalising,

<sup>19</sup>As may be noted, this argument employs itself the informational perspective on logic described in Section 2 in the sense that it relies on the relativization of what can be distinguished to the available conceptual (encoded by Kripke-models) and linguistic resources (modal languages without global modalities).

common logical distinctions guarantee that pooling information will not create further information; the factual information that can be made common for all agents is bounded from above by the distributed information available to all the agents (but see Roelofsen (2006) and Allo (2013*b*, §5) for further qualifications of this claim).

This approach is reminiscent of how Lewis (1969) developed an account of how language could be grounded in convention; a concept he analysed in terms of common belief. My goal here is more modest, and my agreement with Lewis' views is therefore best restricted to the role of common information in successful and reliable interactive practices. In addition to the link with language as convention, conventions and common beliefs are also associated with norms. Pettit (2002), for instance, explicitly incorporates a common belief requirement in his definition of norms as a particular type of regularity (337ff). Given the traditional view that logic is a normative discipline (though spelling out what it is normative for is more complicated), the latter association is surely attractive.



Logical information is information associated with a level of abstraction or the information afforded by the available distinctions at that level of abstraction. In that sense, it is information grounded in what can be distinguished and what can be expressed. The same information can also be understood normatively, for available distinctions also fix standards for valid deduction, and standards for non-triviality. Logical spaces are particularly effective in capturing these two roles, and within this formalism we can often identify the logical information with the total logical space. In an interactive setting, logical information is also common information. Not only is common hard information within a group informationally equivalent to logical information *simpliciter* (i.e. the total logical space), but it is also information that is uniformly and transparently available to all members of the group. Such features underscore the normative role of logic in informational interactions (dialogues, debates, asking and answering questions, ...). Whereas Dutilh Novaes (2015) characterises the normative role of logic by pointing to the strategic advantage of relying on deductively valid inferences in adversary dialogues (the conclusion of a valid argument cannot be challenged if its premises are already accepted), stressing the common nature of logical information gets us closer to cooperative goals of informational interactions by making sure the interaction is reliable. Stenning (2002) explicitly connects the importance of alignment in adversarial situations:

[L]ogic plays its part as a criterion for alignment between interpretations, especially in adversarial communication, when meanings are least clear and passions most inflamed. (Stenning 2002, 6)

Logical information is traditionally conceptualised as a kind of hard information. In this section I made no such assumption, and in the next section I will actively challenge this view by showing that at least some soft information qualifies as logical information.

### 5 SOFT LOGICAL INFORMATION

The qualification of logical information as *knowledge-like* or hard information is a basic, metaphysically neutral, and therefore almost unassailable aspect of our thinking about logical truth and logical consequence. Once we settle for a deductive account of logic, the fit with hard information seems inevitable. My present concern is not with the latter identification, but with the question of whether there is a sensible and useful notion of soft *logical* information that is related to defeasible inference-patterns in the same way as hard logical information is related to deduction. Developing a philosophical account of non-monotonic consequence within an informational conception of logic requires a continuity rather than a sharp cut-off between the deductive and the non-deductive, and a useful account of soft logical information is essential in that respect.

Although the field of non-monotonic logic includes an incredible diversity of formalisms, the central role of preferences in the semantic characterisation of non-monotonic consequence relations has been recognised at least since the seminal work of Shoham (1987). As codifications of common-sense reasoning forms, many non-monotonic consequence relations (and conditionals) are adequately described as formal inference mechanisms that rely on deductive logic to characterise non-deductive inference forms, but only qualify as logics as a matter of family resemblance. Within an informational conception of logic as described in earlier sections, such systems do not require a notion of soft logical information when the preferences they rely on cannot reasonably be considered formal or otherwise “logical”. This much should be uncontroversial: jumping to the conclusion that Tweety can fly upon learning that Tweety is a bird can hardly be described as a logical move.<sup>20</sup> As explained in for instance Velázquez-Quesada (2014), many common forms of ampliative or non-deductive reasoning can be studied as forms of belief-dynamics or dynamics involving soft information. Using the term soft logical information here is neither sensible nor fruitful.

There is, nevertheless, a sub-type of non-monotonic logic that appears to rely

---

<sup>20</sup>This also matches the core concerns of the study of such inferences, which is aimed at the difficulties of managing and combining many (often incompatible) defeasible inferences; as exemplified by the Nixon Diamond.



on a logical kind of soft information (Allo 2016), namely where the relevant inference-patterns that are codified are not supra-classical (and hence could be thought of as truth-preserving in at least a classical sense) and/or where the soft information that underlies these inference-patterns can be described as purely formal instead of factual or material. Paraconsistent logics that compensate the loss of certain classically valid argument-forms by reintroducing them as defeasible inferences are a textbook example, where for instance a defeasible form of the disjunctive syllogism allows one to derive  $B$  from  $\neg A \vee B$  and  $A$  on the assumption that  $A$  behaves consistently. This is the approach taken by Batens (1989; 1999) and Priest (1991). If we only consider the semantic characterisation of the consequence-relations that sanction such defeasible applications of classically valid arguments, we see that they can be based on soft information (i.e. the ordering of models or points in a structure) in exactly the same way as other plausible inferences can be encoded in soft information. Instead of considering states or models where Tweety is a typical flying bird more plausible than those where Tweety is a penguin or an ostrich, they just take consistency as a standard or measure of plausibility.

As before, and adapting results from Allo (2013a), this approach can be captured within a Kripke-style semantics.

**DEFINITION 4 (AB STRUCTURE)** An abnormality structure is a 4-tuple  $\mathfrak{M} = (S, \approx, \preceq, \|\cdot\|_{\mathfrak{M}})$ , with  $S$  a set of points,  $\approx$  an equivalence-relation over  $S$ ,  $\preceq$  a pre-order over  $S$ , and  $\|\cdot\|_{\mathfrak{M}}$  a valuation-functions such that, given a set  $\Omega$  that can be characterised by a possibly restricted logical form, and a function  $Ab : S \mapsto \mathcal{P}(\Omega)$  for which  $Ab(s) = \{\omega \in \Omega \mid s \in \|\omega\|_{\mathfrak{M}}\}$ , we have:<sup>21</sup>

$$s \preceq t \text{ iff } Ab(s) \subseteq Ab(t) \quad (\dagger)$$

Given such a structure (a special case of hard/soft structures), the conditional  $[\overset{A}{\rightarrow}]B$  can (if  $\overset{A}{\rightarrow}$  is defined relative to  $\preceq$  instead of  $\leq$ ) function as a conditional that captures a non-monotonic consequence-relation at the object-level in the same way as a strict implication relates to a monotonic consequence-relation. Specifically, if  $S$  is a logical space that contains all models of the paraconsistent logic  $\mathbf{LP}^{22}$  and  $\Omega$  is the set of all

<sup>21</sup>Note that  $\preceq$  can now no longer be assumed to be (locally) connected.

<sup>22</sup>This variant of strong three-valued Kleene logic is one of the simplest propositional paraconsistent logics. It is obtained by including a third ‘intermediate’ semantic-value and treating both *True* and the intermediate value as designated values. This renders the disjunctive syllogism as well as *modus ponens* and *tollens* invalid, but retains the so-called De Morgan properties. As in Allo (2016, §5), I develop my argument in terms of  $\mathbf{LP}$  and its non-monotonic extensions to retain a certain continuity with existing debates. This focus does not do justice to the wide variety of systems that were developed by Batens and his collaborators. The interested reader should consult Strasser (2014) for a broad but formally highly detailed overview.

atomic contradictions ( $A \wedge \neg A$  for  $A$  atomic), then  $\mathfrak{M}, s \Vdash [\overset{A}{\rightarrow}]B$  will hold whenever  $[s]^\approx \cap \|A\|_{\mathfrak{M}} \neq \emptyset$  and  $B$  can be defeasibly derived from  $A$  using the non-monotonic extension of **LP** driven by the assumption that all inconsistencies, except those that cannot be avoided in view of the information in  $[s]^\approx \cap \|A\|_{\mathfrak{M}}$ , are false.<sup>23</sup>

Abnormality-structures of this kind satisfy both the requirement that the conditional we defined within these structures does not extend classical logic, and that the pre-order  $\preceq$  encodes soft information that can be defined on the basis of a purely formal (or more precisely, *schematic*; see Dutilh Novaes (2011, §2.1.1)) criterion. It is, thus, a *prima facie* example of non-monotonic consequence that is best explained by appealing to a kind of soft logical information. There are, however, some additional wrinkles.



The main worry follows from how Beall (2012) objects to the use of a non-monotonic strengthening of **LP** as an all-purpose logic on the ground that it no longer guarantees that we can never move from true premises to false conclusions. In a reply, Priest (2012) counters that as an ampliative logic, the non-monotonic strengthenings of **LP** were never meant to be truth-preserving in the first place.<sup>24</sup> This isn't particularly helpful, since it merely reiterates the distinction between deductive and ampliative logics, but does not provide additional reasons to treat assumptions of consistency any different from other assumptions of normality or typicality. This is to be expected. As dialetheists, Beall and Priest agree that consistency is not a logical criterion or requirement at all. Beall's subsequent work on the delimitation of logic (Beall 2015; 2014) develops a second argument that has interesting consequences for the distinction between hard and soft (logical) information.

Beall's original objection is this: because the assumption of consistency can lead us from a true **LP**-theory to a false extension of that theory, it violates the requirement of *general reassurance* which he thinks should be a basic requirement on any logic.<sup>25</sup> His second, indirect, objection is based on the claim that "Logic constrains the space of rationally available options" (Beall 2015, 5), but does not tell us which options we should actually retain. In other words, logic tells us about what is

<sup>23</sup>The reader is referred to Allo (2013a) for a more precise statement of this claim, and especially for its more general formulation in Theorems 3 and 5 of that article.

<sup>24</sup>In the original papers, Batens presented similar systems as *corrective* instead of as *ampliative* logics. Classical logic was seen as the correct standard of deduction, but a non-monotonic mechanism had to take care of the cases where it over-generated due to the presence of inconsistencies that the intended standard of deduction couldn't accommodate.

<sup>25</sup>This is a strengthening of the requirement (which is satisfied) that the defeasible assumption of consistency cannot lead us from a non-trivial **LP**-theory to a trivial extension of that theory.

possible (given certain assumptions), but does not add any further structure to this space of possibilities. Although he doesn't use this terminology, he thereby affirms that all logical information is hard information. This is essentially an argument that is based on a presumed division of labour between logical and extra-logical (epistemic, pragmatic) criteria of rationality. The argument from general reassurance, by contrast, is a pure truth-based argument: from a dialetheist standpoint LP must replace classical logic because it is truth-preserving in situations where classical logic isn't (namely in the presence of true contradictions). The requirement of general reassurance implies that a strengthening of LP is only acceptable if it is consistent with the original reasons for adopting LP. Within the conceptual framework developed in the preceding sections, the conclusion of both objections can be avoided.

On an informational conception the requirement of general reassurance is easily dismissed. Since logic-choice is based on an acceptable equilibrium between the discriminatory power (useful distinctions) and deductive strength (useful conclusions) in a given context (and for a particular purpose), triviality can be seen as the total opposite of such an equilibrium: it collapses all distinctions, and only exchanges it for many (arguably useless) conclusions. This can be used to explain what's wrong with the classical reduction of any inconsistent theory to the trivial theory without having to appeal to the truth of at least some inconsistent theories, or even the falsity of the trivial theory. As a consequence, if classical logic is rejected (in a certain context) because it fudges all distinctions (in that context) and replaced by a logic that fares better in that respect, then this logic could also be extended in a non-monotonic way while remaining consistent with the original reasons for abandoning classical logic. Indeed, on this account the requirement of reassurance (see fn. 25) exactly captures the minimal requirement that should be observed.<sup>26</sup>

The argument from a presumed division of labour between the logical (mere possibilities) and the extra-logical (additional structure or criteria) is countered in two moves.



The first move relies on the account of logical information proposed in Section 4, and in particular on the suggestion that logic does not merely supply criteria for

---

<sup>26</sup>One worry may be that I'm at best changing the subject, and indeed have altered logic at its very core by refusing to let considerations of truth to come into play. This can be countered in two ways. First, since too narrow connections between truth-preservation and logical consequence are known to be problematic for a wide range of non-classical logics, an informational conception based on more general informational virtues may in fact better serve our goals. Second, since arguing for the falsity of the trivial theory is perhaps not that much easier than arguing for the falsity of an inconsistent theory (Mortensen 2005, Priest 2006), an informational conception may also allow us to remain on more tractable grounds in that regard.

consequence, consistency, and non-triviality, but also a basis for reliable informational interaction. Soft information can by definition only provide defeasible criteria of consequence, and this feature is shared by common patterns of defeasible inference and the specific “logical” patterns considered in the previous paragraphs. In that sense it may be deemed insufficiently robust to count as logical in any traditional sense. The soft information or preferences on which the latter are based are purely formal preferences (schematically defined, and motivated by a criterion—like consistency—associated with logic), that are defined over the total logical space. Such global preferences thus automatically encode soft *common* information, which supply a common ground and can play a normative role in an interactive setting.

We extend abnormality-structures to the multi-agent case to make this role explicit.

**DEFINITION 5 (MULTI-AB STRUCTURE)** A multi-agent abnormality-structure is a 5-tuple  $\mathfrak{M} = (S, A, \{\approx_i\}_{i \in A}, \preceq, \|\cdot\|_{\mathfrak{M}})$  with  $S$  a set of points,  $A$  a set of agents, each  $\approx_i$  an equivalence-relation over  $S$ ,  $\preceq$  a pre-order over  $S$  that satisfies (+) for a given set of abnormalities  $\Omega$ , and  $\|\cdot\|_{\mathfrak{M}}$  a valuation-function.

Relative to such structures, we can represent the common ( $\approx_A$ ) and individual ( $\approx_i$ ) hard information of the agents as well as their common logical soft information ( $\preceq$ ), and individual soft information ( $\approx_i \cap \preceq$ ). The main difference between factual and logical soft information is that the latter is shared by all agents; it is uniformly and transparently available to all. As such it is similar to hard logical information, and indeed in both cases the result is that given the same hard factual information, two or more agents should have the same individual soft information in virtue of making the same assumptions or using the same logical preferences. This suggests that soft logical information could also be beneficial for cooperative and adversary dialogues, but here things become more complicated. Whereas in the case of hard logical information differences in individual hard information lead to different information but do not allow for misleading interactions, this is no longer the case with soft logical information.

Indeed, it can be shown that the failure of general reassurance exposed by Beall is structurally similar to a well-known feature of belief-dynamics: true information can be used to mislead. Figure 2 below illustrates how the failure of general reassurance can be exploited to mislead without the use of false information. Here, Alice (full circle around the left state) tells Bob, who initially only has the information that either  $p$  or  $r$ , that  $p$  is the case (which is indeed true). In view of this, Bob jumps to the false conclusion that  $\neg r$  (on the implicit, but mistaken assumption that  $p$  behaves consistently).

## Hard and Soft Logical Information




FIGURE 2: No general reassurance

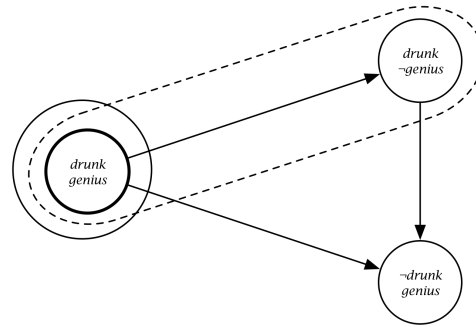


FIGURE 3: "Weinstein" scenario

Figure 4, the "Weinstein scenario" (a standard example used by Baltag & Smets), illustrates a similar phenomenon where Weinstein believes that he's a genius who's definitely not drunk, but will retract his belief that he's a genius upon learning that he's drunk. Here too, an update with true information leads to a false belief because a sensible assumption does not appear to hold.

Consequently, the fact that soft logical information is *common* information is not sufficient to ensure reliable informational interactions, and it can even be used to exploit differences in information and actively convey true but misleading information. This is to be expected, and is entirely in line with how Dutilh Novaes (2015) explains the normative value of deductive arguments in adversarial debates: Only hard information is immune to the threats of information-hiding. When re-described in terms of what can and cannot be distinguished, the result seems equally grim since Bob will rely on distinctions that appear too coarse to properly process the message received from Alice.

But this rendering is misleading, since Bob is actually able to distinguish  $p \wedge \neg p$  from absurdity and only provisionally chooses to ignore this distinction on the basis of currently available information. In other words, the distinctions encoded in soft logical information can act in a self-correcting manner: available distinctions based on hard logical information can be fudged, but can always be retrieved as well. Since soft logical information is associated with the total logical space, it remains in place and acts as a permanent revision-policy. As such, it only ensures reliable informational interaction in the limit, or relative to the total information available to all agents. From the perspective of the individual agent, relying on soft logical information carries risks, but from the perspective of the group or relative to the information in the total system soft logical information can be normative (i.e. it tells you what to take into account to reach certain epistemic goals) in exactly the same

way as hard logical information is, for it encodes how the information-yield should be optimized (in a sense further developed below) once all the available information is made accessible to all agents.



The second move relies on the identification of logic-choice with finding an adequate mix of deductive strength and discriminatory power. That is, a logic is adequate in a given context or for a given purpose if the distinctions it allows ensure that the relevant conclusions can be deduced, and this at a sufficiently fine-grained level. The upshot of appealing to this contrast is that soft logical information can be seen as a further refinement of the balance that can be reached on the basis of hard logical information. If we look at deductive strength, a weaker logic is obtained by removing all (possibly restricted) instances of one or more arguments from a stronger logic. In **LP**, for instance, all instances of the disjunctive syllogism are invalid. If we look at discriminatory power, a more discerning logic will discriminate all instances of some formula's  $A$  or  $B$  (each of a particular, possibly restricted, logical form) that are fudged by a less discerning logic. Again, in **LP** any two contradictions  $A \wedge \neg A$  and  $B \wedge \neg B$  will be logically distinct iff  $A$  and  $B$  can be distinguished as well, whereas classical and intuitionist logic will fudge all of them. This means that, if we only look at hard logical information, the exchange of new distinctions for losses of deductive strength proceeds in a rather crude way. Accommodating a certain inconsistency by distinguishing it from absurdity, and hence blocking the road to triviality, requires one to distinguish all schematically similar inconsistencies and thus giving up a lot more valid arguments than one should have to.

This point can be made even stronger: if in a given context  $p \wedge \neg p$  is or could be derived, then giving up the specific instances of the disjunctive syllogism that could lead to triviality is beneficial, or, equivalently, being able to distinguish  $p \wedge \neg p$  from any other contradiction is equally beneficial. At the same time, rejecting all other instances of the disjunctive syllogism is evidently more costly, and the additional distinctions that are gained through this move are equally irrelevant. In other words, the exchange of inferences for distinctions is hardly optimal if logical distinctions can only be encoded as hard information; a single beneficial exchange is automatically coupled to a much wider range of useless exchanges of potentially valuable inferences for apparently useless distinctions.

Such crudeness is common in logic, and is closely associated to the requirements of truth-preservation and formality. Not only does a focus on truth lead to a very coarse account of semantic content—a notoriously hard problem in relation to the individuation of propositional attitudes like knowledge or belief—, but here

we see that formality enforces distinctions that are too coarse to reach an optimal information-yield (understood as relevant conclusions at a sufficiently fine-grained level). The informational conception of logic is primarily aimed at the first issue (see §2 and Allo and Mares 2012), and extending it with an account of soft logical information allows us to deal with the second issue as well. The cost, in this case, is that we have to give up certain formal features that are traditionally associated with formality. Failures of formality are much less of an issue as one may think. Dynamic epistemic logic, for instance, is a typical example of a family of logics for which uniform substitution fails, and for which the question of *schematic validities* remains an active area of research (van Benthem 2006, Holliday et al. 2013). In that case, schematic validity is closely tied to syntactical form, and the failure of uniform substitution does not pose a threat to the presumed formality of logic. The question we face here is whether the non-uniform treatment of contradictions (some are lumped together, but others are kept apart) or the rejection of instances of arguments is consistent with our previous emphasis on formality (recall that I stressed that soft logical information is based on a formal criterion).

A moderately positive answer is available here as well. The crucial consideration is that a different treatment of several contradictions is only based on the currently available information (the inferential context) and not on the meaning or intended interpretation of the non-logical vocabulary. Whether or not  $p \wedge \neg p$  has to be kept apart from absurdity depends on whether the available information makes  $p \wedge \neg p$  a live possibility,<sup>27</sup> and this depends on what can be derived, but not on what the constituent non-logical terms mean. Since the available information may change over time, a defeasible dynamics is inevitable (what was not initially derivable, may become derivable with new information), but this dynamics is entirely accounted for by the informational context. Hence, the apparent failures of formality are entirely motivated in terms of a broader account of logic-choice, namely the need to optimise information-yield in a specific context and for a particular purpose, and a view of soft information as a refinement of hard information. On such an account, a content or meaning-driven dynamics isn't needed at all.

If we consider what Humberstone calls 1 and 2-synonymy, the contextual nature of distinctions can be made more precise. Here, 2-synonymy coincides with  $\equiv_{\vdash}$  as defined by (Syn $_{\vdash}$ ), whereas 1-synonymy is merely tied to theoremhood in the sense that  $A$  and  $B$  are 1-synonymous iff

$$\vdash C(A) \text{ iff } \vdash C(B).$$

---

<sup>27</sup>What this means depends on the underlying strategy, but it at least depends on whether it can be derived *simpliciter*, or derived as part of a minimal disjunction of inconsistencies.

When applied to nonmonotonic extensions of paraconsistent logics, this means that (*no context*) all contradictions will be 1-synonymous, whereas (*context*) 2-synonymy will vary with the inferential context specified by the premises.

### 6 CONCLUSION

None of the above establishes that certain logics that codify certain defeasible inference forms can or should be adopted as a kind of universal, or general purpose logics. This is intended. Instead, what I believe to have established is that a philosophical account of logical consequence that is grounded in the methodology of the philosophy of information need not be restricted to deductive reasoning-forms, but can be extended to a wider range of informational practices, and ampliative inference in particular. In doing so, a very general account of logic choice based on informational virtues and a deliberately chosen multi-agent perspective have largely replaced or at least generalised the traditional concepts used to characterise logic. This does not lead to a change in subject, but generates a better philosophical account of contemporary logical practices.

### ACKNOWLEDGEMENTS

This paper greatly benefited, at various stages of its development, from the feedback of Marcello D'Agostino, Luciano Floridi and two anonymous referees of this journal.

Patrick Allo  
Marie Skłodowska-Curie Fellow  
*Oxford Internet Institute*  
*University of Oxford*  
patrick.allo@oii.ox.ac.uk  
www.logicandinformation.be

### REFERENCES

- Allo, P. (2013*a*), 'Adaptive logic as a modal logic', *Studia Logica* **101**(5), 933–958.
- Allo, P. (2013*b*), 'The many faces of closure and introspection', *Journal of Philosophical Logic* **42**(1), 91–124.
- Allo, P. (2013*c*), Noisy vs. merely equivocal logics, in K. Tanaka, F. Berto, E. Mares, F. Paoli, S. Rahman and J. Symons, eds, 'Paraconsistency: Logic and Applications',



Vol. 26 of *Logic, Epistemology, and the Unity of Science*, Springer Netherlands, Dordrecht, pp. 57–79.

Allo, P. (2016), ‘Logic, reasoning and revision’, *Theoria* **82**(1), 3–31.

Allo, P. and Mares, E. (2012), ‘Informational semantics as a third alternative?’, *Erkenntnis* **77**(2), 167–185.

Avigad, J. (2015), Mathematics and language, in E. Davis and P. J. Davis, eds, ‘Mathematics, Substance, and Surmise: Views on the Meaning and Ontology of Mathematics’, Springer, pp. 235–55.

Baltag, A. and Smets, S. (2008), A qualitative theory of dynamic interactive belief revision, in G. Bonanno, W. van der Hoek and M. Woolridge, eds, ‘Logic and the Foundations of Decision Theory’, Amsterdam University Press, Amsterdam, pp. 11–58.

Batens, D. (1989), *Dynamic Dialectical Logics*, Philosophia Verlag, München / Hamden / Wien, pp. 187–217.

Batens, D. (1999), *Inconsistency-adaptive logics*, Springer, Heidelberg / New-York, pp. 445–472.

Beall, J. (2012), ‘Why Priest’s reassurance is not reassuring’, *Analysis* **72**(3), 517–525.

Beall, J. (2014), ‘Strict-choice validities: A note on a familiar pluralism’, *Erkenntnis* **79**(2), 301–307.

Beall, J. (2015), ‘Free of detachment: Logic, rationality, and gluts’, *Noûs* **49**(2), 410–423.

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

Blackburn, P., De Rijke, M. and Venema, Y. (2001), *Modal Logic*, Cambridge University Press, Cambridge.

Carnap, R. and Bar-Hillel, Y. (1952), An outline of a theory of semantic information, Technical report, Cambridge, Massachusetts.

Cook, R. T. (2010), ‘Let a thousand flowers bloom: A tour of logical pluralism’, *Philosophy Compass* **5**(6), 492–504.

Cresswell, M., Mares, E. and Rini, A. (2016), *Logical Modalities From Aristotle to Carnap: The Story of Necessity*, Cambridge University Press, Cambridge.

- D'Agostino, M. and Floridi, L. (2009), 'The enduring scandal of deduction', *Synthese* **167**(2), 271–315.
- Dutilh Novaes, C. (2011), 'The different ways in which logic is (said to be) formal', *History and Philosophy of Logic* **32**(4), 303–332.
- Dutilh Novaes, C. (2015), 'A dialogical, multi-agent account of the normativity of logic', *Dialectica* **69**(4), 587–609.
- Etchemendy, J. (1990), *The concept of logical consequence*, Harvard University Press, Cambridge (Mass.).
- Floridi, L. (2008), 'A defence of informational structural realism', *Synthese* **161**(2), 219–253.
- Floridi, L. (2011*a*), 'A defence of constructionism: philosophy as conceptual engineering', *Metaphilosophy* **42**(3), 282–304.
- Floridi, L. (2011*b*), *The Philosophy of Information*, Oxford University Press, Oxford.
- Gabbay, D. (1994), *What is a logical system?*, Studies in Logic and Computation, Oxford University Press, Oxford, pp. 181–215.
- Gabbay, D. and Woods, J. (2003), 'Normative models of rational agency: The theoretical disutility of certain approaches', *Logic Journal of the IGPL* **11**(6), 597–613.
- Gabbay, D. and Woods, J. (2008), 'Resource-origins of nonmonotonicity', *Studia Logica* **88**(1), 85–112.
- Giere, R. N. (2004), 'How models are used to represent reality', *Philosophy of Science* **71**(5), 742–752.
- Glanzberg, M. (2015), Logical consequence and natural language, in O. Hjortland and C. Caret, eds, 'Foundations of Logical Consequence', Oxford University Press, pp. 71–120.
- Gochet, P. and Gribomont, P. (2006), Epistemic logic, in D. M. Gabbay and J. Woods, eds, 'Handbook of the History of Logic Vol. 6', Elsevier, pp. 99–195.
- Hofweber, T. (2014), 'Rayo's the construction of logical space', *Inquiry* **57**(4), 442–454.
- Holliday, W. H., Hoshi, T. and III, T. F. I. (2013), 'Information dynamics and uniform substitution', *Synthese* **190**(1), 31–55.

- Horty, J. F. (2012), *Reasons as defaults*, Oxford University Press, Oxford.
- Humberstone, I. L. (2005), *Logical Discrimination*, Birkhäuser Verlag, Basel, pp. 207–228.
- Jago, M. (2012), ‘The content of deduction’, *Journal of Philosophical Logic* **42**(2), 317–334.
- Leitgeb, H. (2007), ‘Beliefs in conditionals vs. conditional beliefs’, *Topoi* **26**(1), 115–132.
- Lewis, D. (1969), *Convention. A Philosophical Study*, Harvard University Press, Cambridge, Ma.
- MacFarlane, J. (2000), What Does it Mean to Say that Logic is Formal?, PhD thesis, University of Pittsburgh, Pittsburgh.
- Makinson, D. (2005), *Bridges from classical to nonmonotonic logic*, King’s College.
- Mares, E. (2010), ‘The nature of information: a relevant approach’, *Synthese* **175**(0), 111–132.
- Mortensen, C. (2005), ‘It isn’t so, but could it be?’, *Logique et Analyse* **48**(189–192), 351–360.
- Pettit, P. (2002), *Rules, reasons, and norms selected essays*, Oxford University Press ; Clarendon Press, Oxford; New York.
- Priest, G. (1991), ‘Minimally inconsistent LP’, *Studia Logica* **50**(2), 321.
- Priest, G. (2006), ‘Could everything be true?’, *Australasian Journal of Philosophy* **78**(2), June 2000.
- Priest, G. (2012), ‘The sun may not, indeed, rise tomorrow: a reply to Beall’, *Analysis* **72**(4), 739–741.
- Rayo, A. (2012), *The Construction of Logical Space*, Oxford University Press, Oxford.
- Rayo, A. (2014), ‘Reply to critics’, *Inquiry* **57**(4), 498–534.
- Restall, G. (2016), ‘On Priest on nonmonotonic and inductive logic’, *Thought: A Journal of Philosophy* **5**(2), 119–24.
- Rini, A. (2013), The birth of proof: Modality and deductive reasoning, in K. Lodaya, ed., ‘ICLA 2013’, Vol. 7750 of *Lecture Notes in Artificial Intelligence*, pp. 34–49.

- Roelofsen, F. (2006), 'Distributed knowledge', *Journal of Applied Non-classical Logics* **16**(2), 255–273.
- Rumfitt, I. (2015), *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*, Oxford University Press, Oxford.
- Shapiro, S. (2000), *Foundations without foundationalism : a case for second-order logic*, Clarendon Press, Oxford.
- Shapiro, S. (2014), *Varieties of logic.*, Oxford University Press, Oxford.
- Shoham, Y. (1987), Readings in nonmonotonic reasoning, in M. L. Ginsberg, ed., 'A semantical approach to nonmonotonic logics', Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 227–250.
- Smiley, T. (1962), 'The independence of connectives', *The Journal of Symbolic Logic* **27**(4), 426–436.
- Stalnaker, R. (1984), *Inquiry*, MIT Press, Cambridge Ma.
- Stenning, K. (2002), *Seeing reason: Image and language in learning to think*, Oxford University Press.
- Strasser, C. (2014), *Adaptive Logics for Defeasible Reasoning*, Trends in Logic, Springer.
- van Benthem, J. (2006), Open problems in logical dynamics, in D. M. Gabbay, S. S. Goncharov and M. Zakharyashev, eds, 'Mathematical Problems from Applied Logic I: Logics for the XXIst Century', Springer, New York, pp. 137–192.
- van Benthem, J. (2007), 'Dynamic logic for belief revision', *Journal of Applied Non-classical Logics* **17**(2), 129–55.
- van Benthem, J. (2011), *Logical dynamics of information and interaction*, Cambridge University Press, Cambridge.
- Velázquez-Quesada, F. R. (2014), 'Reasoning processes as epistemic dynamics', *Axiomathes* **25**(1), 41–60.