Logique & Analyse **196** (2006), 461–488

LOCAL INFORMATION AND ADAPTIVE CONSEQUENCE

PATRICK ALLO*†

Abstract
In this paper we provide a formal description of what it means to be in a local or partial information-state. Starting from the notion of locality in a relational structure, we define so-called adaptive generated submodels. The latter are then shown to yield an adaptive logic wherein the derivability of $\Box\phi$ is naturally interpreted as a core property of being in a state in which one holds the information that ϕ.

1. *Being Informed and Partial States*

The motivating question of this paper is basically the same as in Floridi (2006), namely: Is there a modal logic in which a formula $\Box\phi$ is naturally interpreted as *holding the information that* ϕ? (or, when subscripted, $\Box_a\phi$ means that a holds the information that ϕ). The position Floridi defends is that within the modal logic KTB[1] the formula $\Box\phi$ exhibits all (and only) the relevant properties of what it means to be informed that ϕ. The argument used to support this claim essentially relies on an examination and interpretation of the relevant axiom-schemes much like that used in the early literature on modal epistemic logic.

The independent argument in favour of a modal adaptive logic for *being informed* presented in this paper is reached in a style which is very different from Floridi's. It departs from it in at least three ways. (i) The system does not straightforwardly rely on an interpretation of *being informed* as truth in all informationally accessible (or indistinguishable) states, (ii) the logic is

*Research Assistant of the Fund for Scientific Research - Flanders (Belgium).
†Junior Research Associate, IEG, Computing Laboratory, Oxford University, UK..

[1] This logic is also known as B, or Br. Semantically, it is characterised by the class of reflexive and symmetric frames; axiomatically it is obtained by adding the axiom-schemes $\Box\phi \to \phi$ and $\phi \to \Box\Diamond\phi$ to the weakest normal modal logic K.

presented in a purely model-theoretic way, and (iii) it allows one to distinguish between the intrinsic and extrinsic properties of *being informed*. We shall nevertheless show that our approach does not contradict the basic tenets of Floridi's axiomatic presentation of KTB as the logic of being informed.

On the formal level, we start from a minimal characterisation — using a generalisation of the frame-semantics for first-degree intuitionistic and relevant logic — of what an information-state is. Hence, we conceive being in an information-state as having a state in an information-ordering as its actual state. Next, it is assumed that one is informed that p if and only if a labelling relation λ assigns p to one's actual state *and* that actual state is a veridical state.[2] We start with the definition of a standard information-ordering, and subsequently give a tentative definition of veridical states.

Definition 1.1: *Let* $(W, \sqsubseteq, *, \lambda)$ *be a labelled structure where* W *is a set of points,* $\sqsubseteq$ *a partial order defined on* W*,* $*$ *a unary operation on* W*, and* λ *a labelling relation assigning standard propositional formulae to points in* W*.* $\sqsubseteq$ *defines an information-ordering on* W *iff: (i) If* $(i, p) \in \lambda$*, then* $(j, p) \in \lambda$ *for all* $i \sqsubseteq j$*. (ii) If* $i \sqsubseteq j$*, then* $j^* \sqsubseteq i^*$*. (iii) If* $i \sqsubseteq j^*$ *then* $j \sqsubseteq i^*$*. (iv) If* $(i, \neg p) \in \lambda$*, then* $(i^*, p) \notin \lambda$*. (v) If* $(i, \phi \wedge \psi) \in \lambda$*, then* $(i, \phi) \in \lambda$ *and* $(i, \psi) \in \lambda$*. (vi) If* $(i, \phi \vee \psi) \in \lambda$*, then* $(i, \phi) \in \lambda$ *or* $(i, \psi) \in \lambda$*.*

The most obvious interpretation of this structure is to consider any set $\lambda(i)$ as a partial state of information. That is, a set of formulae which (a) satisfies an elementary primeness condition (closure under $\vee$ and $\wedge$, see (v) and (vi) above), but (b) is not necessarily complete ($\neg\phi \in \lambda(i)$ iff $\phi \notin \lambda(i)$ is not valid for an arbitrary $i \in W$). Consequently, the relation $\sqsubseteq$ can be seen as a partial ordering on incomplete information-states; its informal interpretation is that of a refinement or informational development. The intended meaning of the $*$-operator is generally considered a more problematic one. Fortunately, in this case, Dunn's interpretation of i^* as i's *maximally incautious twin* (1993) provides us with a simple and useful clue: the operator maps any state onto a state asserting the undenied.[3]

[2] The notion of veridicality can be left largely unanalysed. The only point we have to emphasise, is that veridicality expresses a stronger property than satisfaction or truth in a model. One might, informally, think of it as factually true, true in the actual world, or even — using a terminology we prefer not to rely on here — true in the sense of having a *truthmaker*.

[3] The choice for this framework is one in favour of a very general approach. For instance, negation in this structure behaves like a minimal negation, which is a simple negation, but remains weaker than both intuitionistic (as, for instance, Wright presupposes) and De Morgan negation. The latter two have different but sensible interpretations in terms of partial information-states, and — adopting a pluralist stance with respect to logic and information (see Allo, in press) — it is natural to remain on the common ground between them. Quite

Yet, being a point is this structure is not sufficient to serve as a characterisation of a genuine information-state. Namely, if a set $\lambda(i)$ needs to be consistent, i cannot be any point in W. The labelled structure described above is just too general to enforce the latter, it cannot rule out points to which λ assigns both a formula and its negation. Nevertheless, the combination of the partial order $\sqsubseteq$ and the operation $*$ suffices to discriminate between consistent (no ϕ such that $\{\phi, \neg\phi\} \subseteq \lambda(i)$) and inconsistent points in W. We first expand a bit more on the functioning of $\sqsubseteq$.

Given the constraints (i) and (ii) in definition 1.1, $\sqsubseteq$ enforces a persistence relation between labelled points. Every point is labelled with at least the information its predecessor points are labelled with (1). When combined with the $*$-operator $\sqsubseteq$ equally defines a (symmetric, see (iii) above) compatibility relation between points (2), and trivially also defines consistency as self-compatibility (3). Consequently, information at a point is compatible with the information at all its consistent refinements (4).[4]

$$\text{If } i \sqsubseteq j \quad \text{then} \quad \lambda(i) \subseteq \lambda(j) \tag{1}$$
$$\text{If } i \sqsubseteq j^* \quad \text{then} \quad \lambda(i) \text{ and } \lambda(j) \text{ are compatible} \tag{2}$$
$$\text{If } i \sqsubseteq i^* \quad \text{then} \quad \lambda(i) \text{ is consistent} \tag{3}$$
$$\text{If } i \sqsubseteq j \text{ and } j \sqsubseteq j^* \quad \text{then} \quad \lambda(i) \cup \lambda(j) \text{ is consistent} \tag{4}$$

These insights lead to the obvious conclusion that $\lambda(i)$ is a genuine information-state only if $i \sqsubseteq i^*$.

Additionally, we require $(W, \sqsubseteq, *, \lambda)$ to be such that it contains maximal points. Formally, w is a maximal point in $(W, \sqsubseteq, *, \lambda)$ iff for all ϕ (in a given language), either $(w, \phi) \in \lambda$ or $(w, \neg\phi) \in \lambda$ (but not both).[5] Intuitively, these are points which consistently decide every issue in the same way as possible worlds do. Looking at the clause for negation, it follows that negation only behaves in this way (i.e. classically) on the condition that

differently, minimal negation hardly enforces any interesting property on partial information-states; as Hand (1999) points out in a different context, it does not say anything substantial about the meaning of negation but constitutes a neutral starting point (p. 187). The choice for the Routley-star ($*$) is made for no other reason than its generality and elegance. For instance, adding $i = i^{**}$ suffices to obtain a De Morgan negation. For further generalisations of this framework, see Dunn (1993).

[4] Informally, persistence can be understood as imposing a monotonicity constraint on the ordering such that information at a point is preserved at all its successor points. Compatibility, then, can be seen as a consistency-driven constraint (see also Restall, 1999).

[5] From now on, the letters $i, j, k, \ldots$ will be used to denote arbitrary points in W, $w, w', \ldots$ will only be used when referring explicitly to maximal points in W.

$w = w^*$. Again, this yields a very elegant criterion for maximal points in terms of being their own maximally incautious twin, and an alternative definition of the set of genuine information states as those states which have a maximally consistent refinement $\{i \in W \mid i \sqsubseteq w \;\&\; w = w^*\}$.[6]

Explicitly imposing veridicality within this framework proceeds by defining a designated non-empty subset $\mathsf{WORLD} \subseteq \{i \in W \mid i = i^*\}$ which contains exactly those maximal points in the labelled structure to which λ only assigns veridical information. For the time being, we do not have to be explicit about the number of elements in WORLD; we only require it to be non-empty, and note that if veridicality is a well-behaving concept it should hold that if $w, w' \in \mathsf{WORLD}$, then $\lambda(w) = \lambda(w')$ (where $\lambda(i) = \{\phi \mid (i, \phi) \in \lambda\}$). All other properties of *veridicality* can remain unspecified.

We now define the set of veridical information states $\mathsf{INF\text{-}STAT}$ as those points whose set of maximally consistent refinements is not disjoint from WORLD: $i \in \mathsf{INF\text{-}STAT}$ iff there is a $w \in \mathsf{WORLD}$ such that $i \sqsubseteq w$:

$$\mathsf{INF\text{-}STAT} = \{i \in W \mid \exists w \in \mathsf{WORLD} \;\&\; i \sqsubseteq w\} \tag{5}$$

Evidently, it holds that if $w \in \mathsf{WORLD}$, then $w \in \mathsf{INF\text{-}STAT}$, and that if $i \in \mathsf{INF\text{-}STAT}$, then $i \sqsubseteq i^*$. Consequently w is incompatible with all maximal points not in WORLD. For all non-maximal veridical points, the latter does not hold. In virtue of their partiality, they are compatible with maximal points both in and outside WORLD.

In section 4 more shall be said about the way we define or impose veridicality upon states. At this point, it suffices to note that the sets WORLD and $\mathsf{INF\text{-}STAT}$ have no special status within the logic, nor can the membership of either of these sets be derived from the properties of λ alone. We can now conclude with the formulation of a minimal definition of being informed that p as: the state of an agent which is in an information-state i, where $(i, p) \in \lambda$ and $i \in \mathsf{INF\text{-}STAT}$. This definition shall, further on, be used as both a criterion and a starting point to devise a modal adaptive logic wherein $i \Vdash \Box\phi$ formalises the core properties of being in an information state.

In section 2 we sketch the basics of our modal logic, and define the concept of a generated sub-model. Section 3 modifies this approach by combining it with the preferential models of adaptive logic such as to obtain a localised version of the latter. Finally, in section 4 the topic of what it means to be

[6] This definition is clearly equivalent to the former, but provides a clear-cut connection with the intuition that an information-state is possible (genuine) only if it obtains at a possible world.

in a state of information is taken up again. It is shown to be adequately formalised by the formula $\Box\phi$, and the constraints under which it correctly formalises the stronger relation of *being informed* are outlined.

2. *Locality in Modal Logic*

Introducing what locality in modal logic amounts to, we first have to define a few basics of modal logic. Hereby we almost exclusively rely on the exposition in Blackburn, De Rijke & Venema (2001).

Definition 2.1: (Basic modal language) *Let* PROP *be a set of proposition letters,* $(p, q, r, \ldots)$. *The set* Φ *of well-formed formulae* ϕ *is given by* $\phi ::= p \mid \perp \mid \neg\phi \mid \psi \vee \phi \mid \Diamond\phi$ *for* p *ranging over elements of* PROP.
A dual operator $\Box$ *is defined as* $\Box\phi := \neg\Diamond\neg\phi$*, conjunction and implication are standardly defined in terms of negation and disjunction.*

Definition 2.2: (Frame) *A frame for the basic modal language is a pair* $\mathfrak{F} = (W, R)$ *such that (i)* W *is a non-empty set, and (ii)* R *is a binary relation on* W.

Definition 2.3: (Model) *A model for the basic modal language is a pair* $\mathfrak{M} = (\mathfrak{F}, V)$ *such that* $\mathfrak{F}$ *is a frame and* V *is a map:* $\mathsf{PROP} \to \mathcal{P}(W)$.

Definition 2.4: (Satisfaction) *Let* w *be a state in a model* $\mathfrak{M} = (\mathfrak{F}, V)$*, then we say that a formula* ϕ *is satisfied in that model iff:*
$\mathfrak{M}, w \Vdash p$ *iff* $w \in V(p)$*, for* $p \in$ PROP*,*
$\mathfrak{M}, w \Vdash \perp$ *never,*
$\mathfrak{M}, w \Vdash \neg\phi$ *iff not* $\mathfrak{M}, w \Vdash \phi$*,*
$\mathfrak{M}, w \Vdash \phi \vee \psi$ *iff* $\mathfrak{M}, w \Vdash \phi$ *or* $\mathfrak{M}, w \Vdash \psi$*,*
$\mathfrak{M}, w \Vdash \Diamond\phi$ *iff for some* $v \in W$ *with* Rwv *we have* $\mathfrak{M}, v \Vdash \phi$.

It has become customary to view modal languages essentially as tools for describing relational structures, and more specifically as a tool for modelling the internal perspective on those structures. Blackburn, De Rijke & Venema (2001) describe this core feature of modal logic along the following lines:

> (...) satisfaction is intrinsically *internal* and *local*. We evaluate formulae inside models, at some particular state w. Moreover, $\Diamond$ works locally: the final clause [of the definition of satisfaction] treats $\Diamond$ as an instruction to scan states in search of one where ϕ is satisfied. Crucially, only *successors* of the current state can be scanned by our operators. Much of the characteristic flavour of modal logic

> springs from the perspective on relational structures embodied in the satisfaction definition. (p. 18)

Most of what locality of satisfaction means for a modal language is embodied in the concept of bisimilar models. For the present purpose a less generalised relation between models is of interest, that of generated submodels.

Generated Submodels

To understand the notion of a generated submodel (see Blackburn, De Rijke & Venema, 2001, 55–57), one needs to know what it means for (points in) two structures to be equivalent (given a modal language). We therefore introduce the relation of modal equivalence.

Definition 2.5: (Modal Equivalence) *Let $\mathfrak{M}$ and $\mathfrak{M}'$ be two models of the basic modal language, and let w and w' be two states in $\mathfrak{M}$ and $\mathfrak{M}'$ respectively. Then we say that w and w' are modally equivalent ($w \leftrightsquigarrow w'$) iff they satisfy the same formulae.*

When looking for subsets of a model, the core aim is to find a method for making smaller models out of larger ones without affecting satisfiability. States in a newly constructed model should therefore be modally equivalent to their counterparts in the original model. The most basic method achieving this follows from the notion of a generated submodel.

Definition 2.6: (Generated Submodels) *Let $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}' = (W', R', V')$ be two models. $\mathfrak{M}'$ is a submodel of $\mathfrak{M}$ iff:*
$W' \subseteq W$; $R' = R \cap (W' \times W')$; $V'(p) = V(p) \cap W'$ for all p.
$\mathfrak{M}'$ is a generated submodel of $\mathfrak{M}$ ($\mathfrak{M}' \rightarrowtail \mathfrak{M}$) iff $\mathfrak{M}'$ is a submodel of $\mathfrak{M}$ and if for all w in $\mathfrak{M}'$, if Rwv, then v is in $\mathfrak{M}'$.

Definition 2.7: (Point Generated Submodels) *A submodel of $\mathfrak{M} = (W, R, V)$ generated by the set $X \subseteq W$ is the smallest generated submodel $\mathfrak{M}' = (W', R', V')$ such that $X \in W'$. If X is a singleton set, then it yields a point generated submodel.*

Definition 2.8: (Tree-Structure) *A tree is a relational structure (W, R), where: (i) there is a unique $r \in W$ (the root) such that $\forall w \in W$ it holds that R^*wr, for R^* is the reflexive, transitive closure of R; (ii) for every $w \in W$ distinct from r, there is a unique w' for which Rww'; (iii) R is acyclic.*

Remark that in the above definition it is assumed that R points towards the root, i.e. if Rww' and R is irreflexive, then w' is closer to the root than w.

Proposition 2.1: (Paths in a Structure) *Let (W, R) be a tree-structure, then any $w \in W$ defines a unique path (W', R') in that structure, where $W' = \{w_i \in W \mid R^* w_i w\}$, and $R' = R \cap (W' \times W')$ is a linear order on W'.*

Proposition 2.2: (Finite Submodels) *If a frame $\mathfrak{F}$ is a tree-structure, then any submodel of $\mathfrak{M} = (\mathfrak{F}, V)$ is a finite model.*[7]

An interesting feature of point generated sub-models is their potential use as a mechanism for selecting information up to a certain point (more precisely, it is a model-theoretic characterisation of what a deductive system should select, but not a description of the actual procedure). For instance, in any structure $\mathfrak{F}$ where (i) R is not a universal accessibility relation,[8] and (ii) the model can be unravelled such that a tree-like model is obtained, it holds that for any $w \in W$ the sub-model generated by w is the unique path from w to the root (in the unravelled model).

The illustration below provides an example of this. If the schema represents a tree (or some initial fragment of it), then w is a point in that structure which defines a unique path connecting it to the root (black directed edges). Thus, if formally the path defined by w is the ordered set $\{w, \ldots, r\}$, this corresponds to the path [w $\longrightarrow \ldots \longrightarrow$ root] in the drawing.

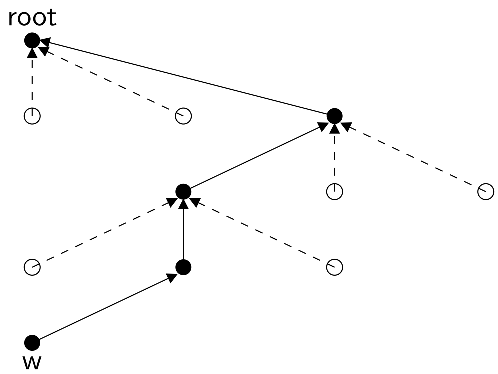

Rww' iff in the figure $w \longrightarrow w'$	R^*ww' iff in the figure $w \longrightarrow \ldots \longrightarrow w'$

[7] A simple example is that of the submodel generated by any n in the well-known structure of the positive integers $(\mathbb{N}, \leq)$. Also clear is that of the submodel generated by 0 in $(\mathbb{Z}, \leq)$ which obviously is not finite as it results in the set of the strictly negative integers.

[8] A relation R is a universal accessibility-relation in (W, R) iff $R = W \times W$, or, equivalently, if R is a reflexive, symmetric, and transitive relation. Informally this means that any point in the structure *sees* any other point in the structure.

If a model (W, R, V) based on a tree (W, R) (as the one in the figure above) is considered, then the point-generated models can be said to contain or select all the formulae (information) assigned (by V) to the points of a unique path. Generally, if we take such a model to be a structured representation of information, a point-generated sub-model is a sensible way (provided that we can give an intuitive, non-formal, interpretation of R) of selecting local information *stored* in or *represented* by that model.

Bi-Modal Case

Extending the basic modal language with a new unary modality allows us to make the language more expressive while keeping additional technicalities limited.[9]

Definition 2.9: (Bi-Modal Language) *Let* PROP *be a set of proposition letters, $(p, q, r, \ldots)$. The set Φ of well-formed formulae ϕ is given by $\phi ::= p \mid \bot \mid \neg\phi \mid \psi \vee \phi \mid \langle P\rangle\phi \mid \langle F\rangle\phi$ for p ranging over elements of* PROP. *Dual operators $[\cdot]$ are defined as $[\cdot]\phi := \neg\langle\cdot\rangle\neg\phi$, for "$\cdot$" being either P or F.*

Definition 2.10: (Defined Modality) *Using both operators, a third* omnitemporal *operator can be defined:*
$\langle A\rangle\phi := \langle F\rangle\phi \vee \langle P\rangle\phi$
$[A]\phi := \neg\langle A\rangle\neg\phi$ *or* $[F]\phi \wedge [P]\phi$

The underlying frame is assumed to remain the same, but the accessibility-relation is now more conveniently represented using $\preceq$, a binary relation defined on W such that $w \preceq w'$ iff $Rw'w$ in the previously used prefix notation. This relation handles accessibility for both basic modalities. The clauses for satisfaction are extended as follows:

Definition 2.11: (Satisfaction) *Let w be a state in a model $\mathfrak{M} = (\mathfrak{F}, V)$ of an underlying frame $\mathfrak{F} = (W, \preceq)$, then we say that a formula $\langle\cdot\rangle\phi$ is satisfied in that model iff:*
$\mathfrak{M}, w \Vdash \langle P\rangle\phi$ *iff for some $v \in W$ with $v \preceq w$ we have $\mathfrak{M}, v \Vdash \phi$.*
$\mathfrak{M}, w \Vdash \langle F\rangle\phi$ *iff for some $v \in W$ with $w \preceq v$ we have $\mathfrak{M}, v \Vdash \phi$.*
$\mathfrak{M}, w \Vdash \langle A\rangle\phi$ *iff for some $v \in W$ with $w \preceq v$ or $v \preceq w$ we have $\mathfrak{M}, v \Vdash \phi$.*

[9] The methods used here are inspired by temporal logics. For this reason the letters P, F, and A are chosen as *labels* for the different modalities and refer to $\langle P\rangle$*ast*, $\langle F\rangle$*uture*, and $\langle A\rangle$*lways*, respectively. The same convention holds for their corresponding boxes $[P]$, $[F]$, and $[A]$ instead of the more traditional H, G (and L) from Prior's systems.

If now $\preceq$ is chosen such that it defines a tree-structure which is branching for successor-points, but connected for predecessor-points, three distinct kinds of point-generated submodels are obtained. The first one for $\langle P \rangle$ was already mentioned for the basic language: it selects the unique path connecting the generator with the root — its determinate past. The second one for $\langle F \rangle$ then selects all paths starting with the generator — its open future, and the third one for $\langle A \rangle$ combines both, but still does not select the structure as a whole — alternative histories and excluded futures remain inaccessible.

Accessibility in a Connected Tree

$\forall i (i \preceq i)$	reflexive
$\forall i \forall j \forall k ((i \preceq j \preceq k) \rightarrow i \preceq k)$	transitive
$\forall i \forall j \forall k ((j \preceq i \wedge k \preceq i) \rightarrow (j \preceq k \vee k \preceq j))$	past-connected
$\forall i \forall j ((i \preceq j \wedge j \preceq i) \rightarrow i = j)$	anti-symmetric

The illustration below represents again a tree-structure. The arrows are chosen such that $w \rightarrow w'$ iff $w' \preceq w$. This means that the arrows represent the accessibility-relation corresponding to $\langle P \rangle$ and $[P]$; the reversed arrows the relation for $\langle F \rangle$ and $[F]$, and finally the symmetric closure of any path that for $\langle A \rangle$ and $[A]$. As in the previous example the black edges represent the generated submodel corresponding to modal equivalence for the operator $\langle A \rangle$.

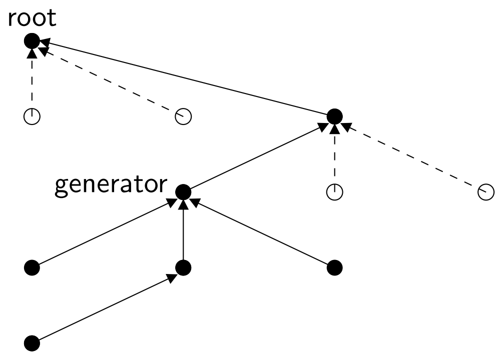

Informally, the features of this structure can straightforwardly be understood as a representation of "loss of possibility", a feature traditionally associated with gaining information:

> In moving to a future time we bypass certain branches. These moves may contain possibilities which are lost as we move past them. Conversely, in moving to a past time new branches become part of our future and possibilities can be gained. Monotonic gain and loss of possibility, with respect to movement into the past and future, is characteristic of future branching time. (Kessler, 1975)

One can, with hindsight, clearly see that the temporal framework used here shares some basic features with the structure defined in section 1. It is, for instance, obvious that the monotonicity referred to above can equally well be represented by either the loss of alternative histories, as by saying that $\sqsubseteq$ represents a persistence relation between states of information.

As known since Hughes (1975), in such structures $\langle F\rangle$ acts as an $\mathsf{S4}$-modality, $\langle P\rangle$ as an $\mathsf{S4.3}$-modality, and $\langle A\rangle$ as a B^{+}-modality. Dismissing the details of respective axioms for these modalities, we only emphasise that the $\langle A\rangle$ operator cannot be interpreted using a transitive-accessibility relation. The reader can easily check that if accessibility for that operator were to be transitive, it would collapse into universal accessibility (as it does for linear time). The logic based on all three modalities will henceforth be referred to as $\mathsf{S4}(.3^{p})$.

Recalling the way in which an information-structure was described in definition 1.1, and combining this with the insight that (i) generated sub-models represent the selection of information up to a point, and (ii) "moving to the future" can be seen as eliminating possibilities, a simple analogy between both can be sketched. So far, the discrepancy between information-structures and rooted tree-like models described by modal logics reduces to the following two points. First, information structures are less constrained, whereas $\preceq$ in the bi-modal logic defines a rooted tree, $\sqsubseteq$ merely enforces a partial ordering. Second, the labelling-function λ (see definitions 1.1) is partial with respect to individual points, the models defined by V (see definition 2.4 and 2.11), on the other hand, are maximally consistent.

The next section introduces the adaptive (generated sub-)models which will be used to bridge the gap between the partiality of λ and the completeness required by models. The discrepancy between partial orders defined by $\sqsubseteq$ and tree-structures is discussed in section 4.

3. *Adaptive Generated Submodels*

Following Batens (2001) an adaptive logic is characterised by its lower limit logic, a set of abnormalities, and a strategy. For this specific case, we opt for the following specification of the standard format:

Lower Limit Logic :	$\mathsf{S4}(.3^p)$
Set of Abnormalities :	$\Omega = \{\langle\cdot\rangle\phi \wedge \langle\cdot\rangle\neg\phi \mid \phi \in \mathsf{PROP}\}$
Adaptive Strategy :	reliability or minimal abnormality

Due to our radical departure from the traditional syntactical use of modal operators in adaptive logics, further modifications to the standard format are required. First, we need to specify that the premises should not be represented as a set of (possibly modal) formulae, but as a labelled graph, or — explicitly using the modal framework — as a partial model:

Definition 3.1: (Structured Premise-Set) *A structured premise-set is a pair $\Gamma^{\preceq} = (\mathfrak{F}, \rho)$ such that $\mathfrak{F}$ is a frame $\mathfrak{F} = (W, \preceq)$ and ρ is a relation: $\rho \subseteq W \times \Phi^{prop}$, for $\Phi^{prop} \subset \Phi$, such that any $\phi \in \Phi^{prop}$ is a non-modal formula of the language defined in 2.9.*

Henceforth, we shall refer to ρ as a set of ordered pairs $\{(w, \phi) \mid w \in W; \phi \in \Phi^{prop}\}$, to $\rho(w)$ as the set of formulae $\{\phi \in \Phi^{prop} \mid (w, \phi) \in \rho\}$ assigned to a certain point w, and to $\rho^{-1}(\phi)$ as the set of points $\{w \in W \mid (w, \phi) \in \rho\}$ to which a certain formula ϕ is assigned.

Using a modal framework and the previously introduced concept of a point-generated submodel, we can construct a localised adaptive semantics. Therefore we first reformulate several notions by giving their point-generated counterpart. To simplify some notation, two new (defined) accessibility-relations will henceforth be used: $i \succeq j := j \preceq i$, and $i \simeq j := i \preceq j \vee j \preceq i$. In the following definitions R is used as a generic placeholder for any of these relations, the corresponding modal operator is written as $\langle\cdot\rangle$ as before (if $\langle\cdot\rangle$ occurs more than once in the same formula, then every occurrence denotes the same operator). As a consequence, most definitions and results listed below should be read as an abbreviation for three distinct definitions, propositions, or theorems. One should also keep in mind that Γ^R, $\Gamma^{\preceq}$, $\Gamma^{\succeq}$, and $\Gamma^{\simeq}$ denote occurrences of one and the same structured premise-set.[10]

Definition 3.2: (Point-Generated Premise-Subset) *Let $\Gamma^R = (W, R, \rho)$ and $\Gamma^R_i = (W_i, R_i, \rho_i)$ be two structured premise-sets, and $i \in W$.*
Γ^R_i is a point-generated premise-subset of Γ^R iff:
$W_i \subseteq W$; $R_i = R \cap (W_i \times W_i)$; $\rho_i = \rho \cap (W_i \times \Phi^{prop})$.
W_i is the smallest subset of W, such that $i \in W_i$ and for all $w \in W_i$, if

[10] The structured premise-set Γ^R could be considered as $(\mathfrak{F}, \rho)$, where $\mathfrak{F} = (W, \preceq, \succeq, \simeq)$ instead of $(W, \preceq)$. When describing generated premise-subsets of $\Gamma^{\preceq}$, $\Gamma^{\succeq}$ and $\Gamma^{\simeq}$ we mean the generated premise-subsets of Γ^R based on $\preceq$, $\succeq$, or $\simeq$ respectively.

wRv, then $v \in W_i$.
$\Gamma_i^{R,\langle\cdot\rangle} := \{\langle\cdot\rangle\phi \mid (w, \phi) \in \rho \,\&\, iRw\}$.

Proposition 3.1: *Γ_i^R and $\Gamma_i^{R,\langle\cdot\rangle}$ contain the same information:*
$\langle\cdot\rangle\phi \in \Gamma_i^{R,\langle\cdot\rangle}$ iff $(w, \phi) \in \rho_i$ and $\Gamma_i^R = (W_i, R_i, \rho_i)$ is a point-generated premise-subset of $\Gamma^R = (W, R, \rho)$.

Point-generated LLL-submodels are directly defined w.r.t. the structured premise-set:

Definition 3.3: (Point-Generated LLL-Submodel) *A point-generated LLL-submodel is a submodel of a premise-set such that, if $\Gamma^R = (W, R, \rho)$, then $\mathfrak{M}_i^R = (W_i, R_i, v_i)$ where $W_i \subseteq W$; $R_i = R \cap (W_i \times W_i)$; W_i is the smallest subset of W, such that $i \in W_i$ and for all $w \in W_i$, if wRv, then $v \in W_i$; v_i is map $\Phi^{prop} \rightarrow \mathcal{P}(W)$, such that if $(w, \phi) \in \rho \cap (W_i \times \Phi^{prop})$, then $w \in v_i(\phi)$.*
And: $v_i(\bot) = \emptyset$; $w \in v_i(\neg\phi)$ iff $w \notin v_i(\phi)$; $w \in v_i(\phi \vee \psi)$ iff $w \in v_i(\phi)$ or $w \in v_i(\phi)$.

$\mathfrak{M}, i \Vdash_{LLL} \langle\cdot\rangle\phi$ iff for some $w \in W$ with iRw we have $w \in v_i(\phi)$).

Proposition 3.2: *A point-generated LLL-submodel $\mathfrak{M}_i^R$ of Γ^R, is an LLL-model of the corresponding point-generated premise-subset Γ_i^R.*

Adaptive models of a premise-set are a specific kind of preferential models, they are obtained by selecting exactly those LLL-models which verify no more abnormalities than strictly necessary according to the premise-set. As the former description does not capture a unique selection of LLL-models, two strategies or ways to select models are described. The generalisation of those selection mechanisms to their point-generated counterparts can be obtained as before:

Definition 3.4: (Disjunctions of Abnormalities — *Dab*) *A Dab-formula Dab (Δ) is the disjunction of the members of a finite $\Delta \subseteq \Omega$. Dab-consequences of Γ_i^R, are the Dab-formulae derivable at i by the LLL from Γ_i^R. A Dab-consequence $Dab(\Delta)$ is minimal iff there is no $\Theta \subset \Delta$ for which $Dab(\Theta)$ is a Dab-consequence.*

Definition 3.5: (Abnormal Part of a Model at a Point) $Ab(\mathfrak{M}, i) = \{\phi \in \Omega \mid \mathfrak{M}, i \Vdash \phi\}$. *The corresponding set of abnormally behaving non-modal formulae is $\{\phi \in \Phi^{prop} \mid \mathfrak{M}, i \Vdash \langle\cdot\rangle\phi \wedge \langle\cdot\rangle\neg\phi\}$.*

Definition 3.6: (Minimal Abnormal Generated Sub-Model) *A point-generated LLL-submodel $\mathfrak{M}_i$ of Γ^R is minimally abnormal iff there is no point-generated LLL-submodel $\mathfrak{M}'_i$ such that $Ab(\mathfrak{M}'_i) \subset Ab(\mathfrak{M}_i)$.*

Definition 3.7: (Unreliable Formulae) *If $Dab(\Delta_1), Dab(\Delta_2), \ldots$ are the minimal Dab-consequences of the point-generated premise-subset Γ_i^R, then $U(\Gamma_i^R) = \Delta_1 \cup \Delta_2 \cup \ldots$ is the set of formulae that are unreliable with respect to Γ_i^R.*

Definition 3.8: (Reliable Generated Sub-Model) *An LLL-submodel $\mathfrak{M}_i$ of Γ^R is reliable iff $Ab(\mathfrak{M}_i) \subseteq U(\Gamma_i^R)$*

Definition 3.9: $\Gamma^R, i \models_{AL^m} \phi$ *iff ϕ is verified in all (relevant) minimally abnormal generated submodels $\mathfrak{M}_i^R$ of Γ^R.*
Or, if we allow for labelled formulae: $\Gamma^R \models_{AL^m} i : \phi$ iff ϕ is true in all (relevant) minimally abnormal generated submodels $\mathfrak{M}_i^R$ of Γ^R.

Definition 3.10: $\Gamma^R, i \models_{AL^r} \phi$ *iff ϕ is verified in all (relevant) reliable generated submodels $\mathfrak{M}_i^R$ of Γ^R*
Or, if we allow for labelled formulae: $\Gamma^R \models_{AL^r} i : \phi$ iff ϕ is true in all (relevant) reliable generated submodels $\mathfrak{M}_i^R$ of Γ^R.

In some cases the set of reliable and minimally abnormal models are identical, when this occurs definitions 3.6 and 3.8 can be replaced by a unique one for the so-called *simple strategy*.

Definition 3.11: (Simple Adaptive Generated Sub-Model) *An LLL-submodel $\mathfrak{M}_i$ of Γ^R is just fine iff $Ab(\mathfrak{M}_i) = Ab(\Gamma_i^R)$*

Standardly this only occurs if the lower limit logic and the set of abnormalities ensure that whenever $\Gamma \models_{LLL} Dab(\Delta)$, there is a $D \in \Omega$ such that $\Gamma \models_{LLL} D$. We shall, however, see in section 4 that a specific restriction on ρ can equally lead to a collapse of both strategies.

A basic but important consequence of the definition of adaptive models is that the lower limit logic together with the set of abnormalities define a so-called upper limit logic (ULL). For instance, if the basic modal logic S5 is chosen as a LLL, and abnormalities are of the form $\Diamond\phi \wedge \Diamond\neg\phi$, where ϕ is a literal, the resulting ULL is the well-known logic Triv, in which every modal distinction collapses (Meheus, 2006). Consequently, a premise set $\Gamma^\Diamond = \{\Diamond\phi \mid \phi \in \Gamma\}$ has ULL-models only if no abnormality is satisfied by its LLL-models. In such a case the adaptive models are exactly the ULL-models.

An analogous definition of ULL-models for a structured premise-set requires, again, some modifications of its standard definition. The basic idea is that we would like to distinguish between ULL-models of the premise-set as a whole, and ULL-models of a point-generated premise-subset. Since we cannot define a single modal premise-set, it has to be noted that ULL-models of the former kind are useless. Next, it is obvious that the plurality of point-generated premise-subsets (due to a choice in both points and accessibility-relation) can result in the existence of ULL-models for Γ_i^R, but not for Γ_j^R, where $\Gamma_i^R \subset \Gamma_j^R$. Likewise $\Gamma_i^{\preceq}$ can have ULL-models while at the same time $\Gamma_i^{\simeq}$ has none.

In view of proposition 3.1 it follows that even if Γ_i^R does not have Triv-models, it still holds that if Γ_i^R has ULL-models, then its modal counterpart $\Gamma_i^{R,\langle\cdot\rangle}$ has matching Triv-models. Yet, as we prefer to stay within a purely model-theoretic approach, the latter will be of little use. When considering the interpretation of $[\cdot]\phi$ as holding the information that ϕ in the next section, we shall see that the existence of ULL-models is highly relevant.

To conclude, we must stress a feature of coupling ULL-models to generated premise-subsets which is easily overlooked. When compared to the standard conclusion that adding the axiom-scheme $\Diamond\phi \rightarrow \Box\phi$ to S5 unambiguously results in Triv, the present approach does not result in a total collapse of the modal structure. Concretely, the systematic reduction of standard notions into their localised counterparts blocks, so to speak, the standard move from several worlds to one world. Whereas in the move from S5 to Triv every modal distinction is lost, the present approach collapses single modalities, but leaves some iterated modalities unaltered.[11]

When presented model-theoretically, showing that an adaptive logic is well-behaving at least amounts to proving the property of strong reassurance. This property ensures that the method used for selecting LLL-models does not yield infinite sequences of less and less abnormal models.

Theorem 1: (Strong Reassurance) *If $\mathfrak{M}_i$ is a point-generated LLL-submodel which is not a minimally abnormal / reliable model, then there is a minimally abnormal / reliable point-generated LLL-submodel $\mathfrak{M}_i'$ such that $Ab(\mathfrak{M}_i') \subset Ab(\mathfrak{M}_i)$.*

[11] It is left as an exercise to check that combined modalities like $\langle P\rangle[F]\phi$ are such unaffected modalities, since, even if true at i there is not necessarily a j such that iRj and ϕ is true at j.

We prove this theorem indirectly by showing that the results of the proof based on the standard format[12] equally apply for this logic. The basic idea is that every structured premise-set, premise-subset, model and submodel can be transformed into a set of modal formulae true at some point in the structure (see proposition 3.1). As the latter formulation matches the standard format, we show that strong reassurance also holds for the former. Remark that this translation is only effective because the abnormal part of a model at a point (definition 3.5) was already specified in terms of formulae satisfied by generated sub-models, and not in terms of the generated submodels themselves (for this specific problem, see Batens, 2000).

Proof. Let $\mathfrak{M}_i^{\langle\cdot\rangle}$ be the set $\{\langle\cdot\rangle\phi \mid \mathfrak{M}, i \Vdash \langle\cdot\rangle\phi\}$. Then from propositions 3.1 and 3.2: (i) $\mathfrak{M}_i$ is a point-generated LLL-submodel of Γ^R iff $\mathfrak{M}_i^{\langle\cdot\rangle}$ is an LLL-model of the point-generated premise subset $\Gamma_i^{R\langle\cdot\rangle}$. (ii) $\mathfrak{M}_i$ is a minimally abnormal / reliable point-generated LLL-submodel of Γ^R iff $\mathfrak{M}_i^{\langle\cdot\rangle}$ is a minimally abnormal / reliable LLL-model of the point-generated premise subset $\Gamma_i^{R\langle\cdot\rangle}$. And by the standard format: (iii) If $\mathfrak{M}_i^{\langle\cdot\rangle}$ is a point-generated LLL-submodel which is not a minimally abnormal / reliable model, then there is a minimally abnormal / reliable point-generated LLL-submodel $\mathfrak{M}_i'^{\langle\cdot\rangle}$ such that $Ab(\mathfrak{M}_i'^{\langle\cdot\rangle}) \subset Ab(\mathfrak{M}_i^{\langle\cdot\rangle})$. □

4. *The State of Being Informed in Dynamic Perspective*

Although the most straightforward interpretation of what this logical system does, is the modelling or reconstruction of the consistency-driven process of acceptance and rejection of data (Allo, 2005), it equally yields a modal operator with a non-standard meaning which is interesting in it own right. Namely, $[\cdot]\phi$ is true at a point iff ϕ is asserted at an accessible point, and remains undenied at any other accessible point (undenied receives here the *strong* interpretation of undenied by any combination of accessible points).[13] Hence, this modality can be understood as a realistic interpretation for, among others, the relation *holding ϕ to be information* (even if there is only evidence

[12] The standard format is a general characterisation to which adaptive logics can comply. If an adaptive logic fits into the standard format, it inherits all the properties which can be proved on the basis of the standard format alone (Batens, 2001, in press).

[13] Keep in mind that being undenied is, strictly speaking, only a correct interpretation when assuming the simple strategy (def. 3.11). For the reliability strategy a disjunction of abnormalities like $(\Diamond\phi \wedge \Diamond\neg\phi) \vee (\Diamond\psi \wedge \Diamond\neg\psi)$ suffices, if minimal, as a *denial* of ϕ.

for treating ϕ as reliable — in the sense of not contradicted — data).[14] The adaptive consequence relation defined by the preferential models introduced in the previous section allows one to select or derive from an input of data exactly those parts and consequences of the incoming data of which it can safely (i.e. coherently) be assumed that they can be treated as if they were genuine pieces of information. For instance in the following structure (for the sake of simplicity we use a linear structure), $[\cdot]\phi$ holds at the rightmost state only if it accesses no state (or set of states) from which the abnormality of ϕ is (jointly) derivable.

$$\phi \longleftarrow \psi \longleftarrow [\cdot]\phi \qquad \cdots$$

However the stronger (and equally important) relation of *holding the information that* ϕ (in the sense of being informed that ϕ only if ϕ is true) is not a *prima facie* interpretation of $[\cdot]\phi$, it does — when applied to the proper context — emerge as a sound (though maybe incomplete)[15] interpretation of the former relation. Effectively obtaining this context requires us to impose some very stringent constraints upon the structured premise-set we start from.

The basic idea is to construct the relation ρ (which determines a structured premise-set $\Gamma^R = (W, \preceq, \rho)$) such that a well-chosen subset of points in the structure only gets assigned true formulae (in the sense of veridical information, not in the sense of logically valid formulae). This should limit the scope of nonmonotonicity within the logic in a way comparable to the enforcement of a monotonicity-constraint by $\sqsubseteq$ in an information-ordering. The formulation of this stronger interpretation of $[\cdot]\phi$, and the necessary restrictions on the structured premise-set to obtain it, constitute the core aim of this section. Two aspects need to be dealt with: the nature of the structure (the partial order induced by $\sqsubseteq$ vs. the tree-structure enforced by $\preceq$), and the compatibility and veridicality of the data ρ assigns to points in the structure.

A first step is to modify ρ such that $(W, \preceq, \rho)$ can be completed into an information ordering $(W, \sqsubseteq, *, \lambda)$ (where both structures are identical), and where the adaptive models of $\Gamma^{\preceq} = (W, \preceq, \rho)$ are themselves models of the information-ordering $(W, \sqsubseteq, *, \lambda)$.

[14] The relation expressed by *holding* ϕ *to be information* has a normative strength comparable to that of *strong belief* in a combined logic of knowledge and belief containing the axiom $B\phi \rightarrow BK\phi$ (see Stalnaker, 2006, sect. 3).

[15] Again, with a reference of to the notion of strong belief, one could wonder if true, strong belief is sufficient for knowledge, and by analogy doubt whether *holding* ϕ *to be information*, and ϕ *is true* are jointly sufficient to *hold the information that* ϕ.

We start with a note on the requirement that $(W, \preceq)$ and $(W, \sqsubseteq)$ ought to be the "same" structure. The reason why this is a prerequisite is obvious. If ρ should be completable into λ, the structural properties of $(W, \preceq, \rho)$ should equally be retained in $(W, \sqsubseteq, *, \lambda)$. Consequently, we need the same set and the same ordering of points. Seeing why this requirement is harmless is relatively straightforward too. The basic insight is this: there exists no structure $(W, \sqsubseteq)$ which is prior to the tree-structure $(W, \preceq)$, definition 1.1 merely lists the minimal constraints we put on an information structure. In other words, the information-structure based on $(W, \preceq, \rho)$ will rely on a more constrained $\sqsubseteq$ than is strictly required to be an information-structure: it will be a tree-structure, and hence also a partial ordering. Put simply, not every information ordering will be the completion of a structured premise-set based on a tree-structure, but every such tree-structure will be completable into an information-structure if ρ satisfies the further conditions outlined below. However, it does hold that every unravelled information-structure will be the completion of a structured premise-set with the required ρ; and this even when $(W, \sqsubseteq, *, \lambda)$ contains inconsistent points.[16]

Noting, as a final remark on the structure itself, that every information-structure *can* be unravelled, i.e. transformed into a tree-structure without loss of factual nor structural information,[17] our choice to restrict our attention to tree-structures is definitely shown to be unobjectionable.

Taking an unravelled information-structure as the structured premise-set $\Gamma^{\preceq}$ we wish to complete with an adaptive consequence-relation — though trivially being a perfectly sound strategy — reduces the role of the adaptive consequence relation to an absolute minimum. We therefore opt for a weaker characterisation of the structured premise-set we need, namely one in which ρ does not assign the information one effectively holds at a point (like λ does), but only the information one acquires at that point. Points in the structure $(W, \preceq, \rho)$ primarily refer to the input, not to the information-states themselves.

[16] This might sound rather unintuitive, and we shall have to come back to this issue later. For now, it suffices to note that the consistency requirement only occurs in definition 3.3, that is at the level of (sub-)models and not at the level of structured premise-sets.

[17] The formal result on which this argument is based, is the following. Given a modal language containing only diamonds, and a rooted model $\mathfrak{M}$, there is a tree-like model $\mathfrak{M}'$ such that every formula of that language which is satisfiable in the rooted model, is also satisfiable in the tree-model. See: Blackburn, De Rijke & Venema (2001): proposition 2.15, pp. 62–63.

Provisionally, we require that any structured premise-set (to which we sometimes refer as either a labelled tree or a partial model) constructed using the newly defined ρ (to which we shall refer as ρ_g) should be compatible with a λ from an information-ordering based on the same tree-structure in the sense that the former can be completed into the latter. In other words, the following inclusion should hold: $\rho_g \subseteq \lambda$ for ρ_g assigns to any point the newly acquired (but possibly redundant) information, and λ defines a standard information ordering. As an example, if — for a simple linear case — a standard information ordering is represented by a sequence of points / sets of formulae $i \sqsubseteq j \sqsubseteq k$, resp. $\lambda(i) \subseteq \lambda(j) \subseteq \lambda(k)$, a corresponding sequence of points / sets of newly acquired (non-redundant) information can be devised as $i \preceq j \preceq k$, resp. $\lambda(i)$; $(\lambda(j) \setminus \lambda(i))$; $(\lambda(k) \setminus \lambda(j))$ (obviously, the set-inclusion does not hold any more).

Before we move on to the topic of veridicality of input, a final word on inconsistent points remains to be said. Consider the following setting: let $(W, \sqsubseteq, *, \lambda)$ be an unravelled information-ordering containing some inconsistent points. As mentioned above, even in that case a corresponding structured premise-set $(W, \preceq, \rho)$ can be constructed.[18] Using the point-generated sub-models defined in the previous section, this fact gives rise to two major problems. A first one is the possible lack of point-generated ULL-models — an issue to which we have to pay attention, but which eventually turns out to be harmless; a second one is the possible lack of point-generated LLL-models, which forces us to exclude some seemingly admissible structured premise-sets. Whereas the first problem arises as soon as a point-generated premise-subset validates an abnormality (if ρ follows λ's closure under $\vee$ and $\wedge$ this necessarily is a simple abnormality, not a disjunction of abnormalities), the second one can only arise when a single point validates a contradiction ($\rho(i) = \perp$ for some i), informally, the inconsistency of a single input.

With regard to the second problem, two strategies can be adopted. The most general and coherent approach to the problem, is to characterise the LLL-models themselves in terms of an inconsistency-adaptive logic (see the combined adaptive logic in Meheus (in press)). However, for reasons of focus, it is preferable to explicitly exclude premise-sets containing inconsistent points, and thereby fix the intended domain of application of the adaptive logic to that of consistent input at any point (see e.g. Meheus (2006)). The reasonability of this approach is warranted in two ways: it is a standard assumption in adaptive logics that the LLL determines the intended domain of

[18] Note that an unravelled information-ordering does not have a unique corresponding structured premise-subset, unless we require the latter to be non-redundant.

application, and within this restricted domain of application both the combined and simple adaptive logic validate exactly the same consequences.

For the first problem, then, two possible cases shall be considered separately (see (ii) and (iii) in proposition 4.1). Generally, this section is in the first place concerned with genuine information states. These are, as we shall see, states where only abnormalities of the form $\langle F\rangle\phi \wedge \langle F\rangle\neg\phi$ can occur. In turn, for those cases where (equally harmless) abnormalities of the form $\langle P\rangle\phi \wedge \langle P\rangle\neg\phi$ arise, it suffices to mention that the existence of LLL-models warrants the non-triviality for all formulae, and the selection of LLL-models in the adaptive logic warrants the maximisation of the boxed formulae in those cases too.

As we did for the basic information orderings, we now have to define a set of veridical points in the structured premise-set based on ρ_g. Therefore a modified set WORLD_g is defined as the set of maximal paths in the structure to which ρ_g exclusively assigns veridical information. The set of veridical points $\mathsf{INF\text{-}STAT}_g$ is then the set of those points which lie on a path that is in WORLD_g. The intended relation between the information-ordering, the structured premise-set, and their matching sets WORLD and $\mathsf{INF\text{-}STAT}$ is (independently of the existence of inconsistent points) captured by the following constraints:

Constraint 4.1: (λ and ρ_g define the same set of actual worlds) *Let* $(W, \preceq, \rho_g)$ *and* $(W, \sqsubseteq, *, \lambda)$ *be two partial models or labelled trees which only differ w.r.t. to* ρ_g/λ*. Then for every* $i \in W$*, we say that* $i \in \mathsf{WORLD}$ *w.r.t.* λ *iff* $(\{j \mid j \preceq i\}, \preceq) \in \mathsf{WORLD}_g$*, i.e. actual worlds in* WORLD_g *are paths, see constraint 4.3*

Constraint 4.2: (λ and ρ_g define the same set of veridical states) *Let* $(W, \preceq, \rho_g)$ *and* $(W, \sqsubseteq, *, \lambda)$ *be two partial models or labelled trees which only differ w.r.t. to* ρ_g/λ*. Then for every* $i \in W$*, we say that* $i \in \mathsf{INF\text{-}STAT}$ *iff* $i \in \mathsf{INF\text{-}STAT}_g$*.*

Constraint 4.3: (Maximal paths in $(W, \preceq, \rho_g)$ define possible worlds) *Let* $(W, \preceq, \rho_g)$ *and* $(W, \sqsubseteq, *, \lambda)$ *be two partial models or labelled trees which only differ w.r.t. to* ρ_g/λ*. Then for every maximal point* w *in* $(W, \sqsubseteq, *, \lambda)$*, the following inclusion should hold:* $\{\phi \mid j \preceq w \;\&\; (j, \phi) \in \rho_g\} \subseteq \{\phi \mid (w, \phi) \in \lambda \;\&\; w = w^*\}$.

Remark: given the failure of $i \sqsubseteq i^*$ for arbitrary i in $(W, \sqsubseteq, *, \lambda)$ it does not hold that every path in $(W, \preceq, \rho_g)$ is the initial fragment of a maximal path matching a possible world. Additionally, even if $i \sqsubseteq i^*$, i may have an inconsistent refinement j, and for the same reason as before, even if i defines the initial fragment of a maximal path in $(W, \preceq, \rho_g)$, it does not follow that it is the initial fragment of maximal paths only.

Remember that in the original definition of veridicality in an information-structure, the set WORLD was merely required to be non-empty. If, however, we only consider unravelled information-structures, it is preferable to explicitly exclude the possibility of WORLD being a singleton. Even more, if we want the process of receiving data to be non-deterministic, no set $\mathsf{i} - \mathsf{WORLD} = \{w \in \mathsf{WORLD} \mid i \in \mathsf{INF\text{-}STAT}\ \&\ i \sqsubseteq w\}$ should be a singleton. By the same token, every point in $\mathsf{INF\text{-}STAT}_g$ should lie on more than one veridical path in WORLD_g; that is, at any point there should be more than one way to receive data *and* remain in a veridical state. More generally, every point should *see* more than one identical possible world, or *lie* on more than one path to which ρ_g assigns exactly the same data.

We now can proceed to the interpretation of a formula $[\cdot]\phi$ satisfied at a point in all adaptive point-generated sub-models of $\Gamma^{\preceq} = (W, \preceq, \rho_g)$. Which means for the different modalities: $i \Vdash [P]\phi$ iff i has at least one predecessor-point $j \Vdash \phi$, and no subset of predecessor-points which jointly deny ϕ (again, relative to what the latter *means* in the adaptive strategy used). $i \Vdash [F]\phi$ iff i has at least one successor-point $j \Vdash \phi$, and no finite subset of successor-points which jointly deny ϕ. $i \Vdash [A]\phi$ iff i has at least one accessible point $j \Vdash \phi$, and no finite subset of accessible points which jointly deny ϕ, where "joint denial" is systematically understood as being closed under LLL-consequence.

The basic interaction between the adaptive consequence relation, and the existence of an information-ordering based on the structured premise-set, is governed by the following proposition.

Proposition 4.1: *Let $(W, \preceq, \rho_g)$ and $(W, \sqsubseteq, *, \lambda)$ be two partial models or labelled trees such that $i \preceq j$ iff $i \sqsubseteq j$. If in accordance with constraints 4.1–4.3 it holds that $\rho_g \subseteq \lambda$ and λ defines an (unravelled) information-ordering in the sense of definition 1.1, then:*
(i) If $j \sqsubseteq j^*$ holds for all $i \sqsubseteq j$ in $(W, \sqsubseteq, *, \lambda)$: if $\phi \in \lambda(i)$, then $i \Vdash [A]\phi$. That is, no abnormality is derivable at i in $(W, \preceq, \rho_g)$ if i is a persistently consistent point in $(W, \sqsubseteq, *, \lambda)$. Hence every adaptive sub-model generated by i is an ULL-model.
(ii) If $i \sqsubseteq i^*$ holds in $(W, \sqsubseteq, *, \lambda)$: if $\phi \in \lambda(i)$, then $i \Vdash [P]\phi$. That is, no $\langle P\rangle$-based abnormality is derivable at i in $(W, \preceq, \rho_g)$ if i is a consistent point in $(W, \sqsubseteq, *, \lambda)$. Hence every adaptive sub-model generated by i w.r.t. $\preceq$ is an ULL-model.
(iii) If $i \not\sqsubseteq i^*$ holds in $(W, \sqsubseteq, *, \lambda)$: no adaptive models generated by i are ULL-models, hence the inference between $\phi \in \lambda(i)$ and $i \Vdash [\cdot]\phi$ does not hold in general. Yet, if ρ satisfies the same primeness conditions as λ, the

reliable and minimal abnormal models collapse. Hence, every point closed under $\vee$ *and* $\wedge$ *has simple adaptive models.*

As we are essentially after a characterisation of the minimal context in which $[A]\phi$ formalises the notion of being informed, we can consistently restrict our attention to the first two cases (and especially (ii) since it describes a genuine information-state *and* includes (i) as a special case). Given the nature of ρ_g and its relation to λ the proviso that $[\cdot]\phi$ only holds in case the denial of ϕ is not accessible is vacuous except in the case of successor-points. Concretely, an abnormality $i \Vdash \langle\cdot\rangle\phi \wedge \langle\cdot\rangle\neg\phi$ can only occur if $\langle\cdot\rangle$ is $\langle F\rangle$. If, however, $\langle\cdot\rangle$ were $\langle P\rangle$ the occurrence of an abnormality would contradict the original assumption that the underlying premise-set can be extended to an information-ordering in which i is a genuine information-state.

Despite the limited effect of the adaptive approach in any genuine information-state, the only correct formalisation of being in a state in which one holds the information that p, is that of $i \in \mathsf{INF\text{-}STAT}_g$ (which satisfies all properties described in (ii) above), and $i \Vdash [A]p$, for it is the only formulation which captures all three of the following properties: [P-safe] having received the undenied information that p (a property of predecessor-points), [F-safe] trusting that p is undeniable (a purported property of successor-points), and [VER] p is true (a property of the actual state).

A possible objection to this formal characterisation of being in an information state in which one holds the information that p, is that being in a state $i \in \mathsf{INF\text{-}STAT}$ and $i \Vdash p$ ensures on its own the joint satisfaction of [P-safe], and [VER], and indirectly that of [F-safe].[19] If all that is required to capture the notion of *being informed* lies in the definition of semantic information as veridical well-formed meaningful data, then clearly being true of a veridical state does the trick and there is no need for an adaptive logic to model its properties. The reason why we hold that the sheer fact of being in a veridical information state does not capture all relevant features of *being informed* is tied to the correct assessment of how ρ_g/λ, WORLD, and $\mathsf{INF\text{-}STAT}$ interact; or more importantly, fail to interact in an important sense.

Veridicality of Points, and Veridicality in the Structure

The problem we face is the following. All three conditions listed above are necessary, but being in a veridical state apparently provides a sufficient condition on its own. Nevertheless, it has to be stressed that being a veridical

[19] Namely, if i is a genuine information state, it could still have refinements which are themselves not genuine information states (see case (ii) in proposition 4.1) Informally, while i might have refinements or successors which deny the information one holds in i, such states could never be one's actual state.

state is not a property of the state itself in the sense that i is a veridical point iff there is a χ such that $i \Vdash \chi$, and "χ" is true iff $i \in \mathsf{INF\text{-}STAT}_{(g)}$. It is only a property in the weaker sense that $i \in \mathsf{INF\text{-}STAT}_{(g)}$. Hence, the information that one is in a veridical state — a property denoted by the hypothetical formula χ — is not available within that state. The only *meta-information* one might be tempted to accept as being contained within the state is that $i \Vdash [A]\phi$ expresses the undefeated assumption that i has a maximal successor-point w which can be understood as a possible world and makes ϕ true (note the analogy with $i \sqsubseteq w$ and $w = w^*$). That it *cannot* be defeated, on the other hand, is only ensured in virtue of the properties of a genuine information-state. Namely, [P-safe] is ensured by $\rho_g \subseteq \lambda$ and $i \sqsubseteq i^*$ (itself a trivial consequence of $i \sqsubseteq w$), and [F-safe] by the fact that (given $i \Vdash [P]\phi$, or equivalently $\phi \in \lambda(i)$) no abnormality of the form $\langle F\rangle\phi \wedge \langle F\rangle\neg\phi$ can be derived at i on the basis of genuine information-states alone.

Unfortunately, the conditions for $[A]$ — even if indefeasible by any genuine information-state — do not ensure that i has a maximal successor-point which defines the actual world. To better grasp the latter fact, one should reconsider that case (ii) of proposition 4.1 holds for all consistent points in the structure, not just for the veridical ones.

In answering this objection, two options remain open. Either one rejects the analysis of $[A]p$ as holding the information that p on the ground that it fails to really enforce the veridicality requirement (and only weakly endorses a coherence requirement); or, one can argue that its failure to enforce veridicality as an intrinsic property of an information-state is not a drawback or weakness of the present analysis, but a perfectly sound consequence of what it means to hold information. We shall, in other words, argue that a correct analysis of veridical information is coherent with the idea that truth does not supervene on the concept of information (a point made by Floridi, 2005), while it cannot exclude that truth does supervene on the description of the underlying internal states.[20] Namely, even if $i \in \mathsf{INF\text{-}STAT}_{(g)}$, there is no χ such that $i \Vdash \chi$ can, for an agent whose actual state is i, contain the information that $i \in \mathsf{INF\text{-}STAT}_{(g)}$.[21]

[20] Compare with the possibility of falsely believing that one knows that p — itself an important feature of strong belief.

[21] One could, however, assign unique names out of a special set of propositional parameters or nominals — unique names for states, as used in hybrid modal logic (see Blackburn, De Rijke & Venema, 2001, 7.3) — to all the points in W. Obviously, there would be a well-defined subset of nominals which are names of veridical states, but no agent could be informed of either the name of its actual state, nor the fact that it is a name of a veridical state. This, clearly ineffective method, is the closest one could get to the hypothetical formula χ.

Dealing with the first half of this dilemma can be done rather swiftly, for one simply has to recall that the analysis we gave of what it means to hold the information that p. That is, affirming that it necessarily relies on two distinct features: an internal one expressed by $i \Vdash [A]p$, and an external one expressed by $i \in \mathsf{INF\text{-}STAT}_{(g)}$. Never in the course of the present analysis was it claimed that the conditions [P-safe] and [F-safe], both relying on the operational meaning of $[A]$, could constitute a criterion for veridicality. Hence, it would be unfair to ask more of internal-states than what they were meant for in the first place.

The second part of the answer requires a more elaborate approach, and encompasses a defence of two theses: (i) an internal state which satisfies its own veridicality fails to be internal, and (ii) internal states do matter.

Support for the first thesis follows from the nature of an information-state, and even more from our understanding of what factual information is (and does). Namely, an information-state is nothing more than a set of statements, sentences, or propositions (on their precise nature we can remain undecided), possibly closed under some consequence relation. Factual information is well-formed meaningful data which additionally happens to be true. However, following the traditional Dretskian analysis, false and true messages are not different *qua* messages — they both constitute semantic content. Yet, only a true message can inform me of its content.

> (...) it makes no difference if one person knows that the signal he is receiving is reliable and the other does not. As long as the signal is reliable, whether or not it is known to be reliable (Dretske, 1999, 81).

In that sense it seems obvious to deny internal states to carry their veridicality as an intrinsic property, for they can be considered as merely (logically closed) sets of messages. But this can exactly be considered as an argument for rejecting the internal analysis altogether. Which brings us to the second thesis: affirming the relevance of a local or internal analysis of information-states themselves.

The main reason to consider information-states as internal states, is that they provide a coherent interpretation of partiality. When we hold information, we rarely hold all veridical information, and the ignorance which unavoidably comes with this partiality, is elegantly rendered when considered from a local perspective. Still, the mutual coherence of partiality and locality does not constitute a conclusive argument for the relevance of internal states, and certainly not for states which fail to support their veridicality by themselves. The usefulness of models which adopt a purely external perspective does, on the other hand, not constitute an argument against the soundness of our approach either. On the positive side, we claim that looking at internal states provides a better grip on what it means to hold information. Moreover, we hold that such insights are generally not available to the more traditional doxastic and epistemic models. Making this an effective claim can only be achieved by considering not only the states themselves, but also the fact that holding the information that p, requires a state i such that $i \Vdash [A]\phi$ describes — or is true of — an agent's actual state.

Considering those agents in the most general fashion, there is no reason to assume them to be any more powerful than what is strictly required for holding syntactical strings and maybe, but not necessarily, manipulate these strings. Agents which are in a particular information state, or even databases which simply *are* information-states generally do not need an internal theory of truth. It may of course be objected that when assigning formulae to an information-state, a satisfaction-relation is implicitly assumed. But, in the same sense as a model-theory is not a semantics — let alone a theory of meaning — we hold that a satisfaction-relation alone does not constitute a theory of truth capable of yielding a full-fledged account of veridicality.[22] What we aim at, is that while the internal language of such an entity should not contain an internal truth-predicate, the lack of a concept of truth does not preclude that its actual state either is or is not a veridical state. This is precisely the intended meaning of the thesis that truth does not supervene on our analysis of information, but that it does supervene on an agent's internal state (it is not part of the information available to the agent).

Evidently, such claims have important repercussions on how we understand the concept of truth within a theory of (semantic) information. Actually, two distinct reasons to deny internal states a theory of truth of their own come into play. The first reason is a straightforward consequence of the underlying realism of information theory when interpreted in a mainstream Dretskian fashion. If semantic content can be said to qualify as genuine semantic information independently from the existence of an informed agent, then obviously it becomes very hard, if not plainly incoherent, to try to give an account of veridicality based on the internal states of an agent the theory does not presuppose in the first place. But, giving an account of internal states is in itself not incoherent nor superfluous. It tells a part of the story the standard theory cannot explain, and, when refraining from making truth-claims, it does not contradict the Dretskian picture. This constitutes a first — strong — ground to reject an analysis of veridicality based on internal states. However strong on the pain of being incoherent, the previous argument is far from being absolute as it appears to fail when considered outside the scope of informational realism.

[22] In this context, it is useful to recall that logic as such cannot account for the *badness* of falsity (Hand, 1999) — true and false receive a perfectly symmetric treatment. When it comes to information, only veridicality matters.

The second argument provides a weaker, but realism-independent, ground for developing a truth-free account of internal states. Once the assumption that information is informee-independent is left out of consideration, there is room for an analysis of veridicality that does not essentially and exclusively rely on realist assumptions. Not only does truth-talk within internal states stop from being incoherent when departing from the externalist picture, an actual analysis of internal truth-talk seems to arise from our formal analysis.

The latter being an extensive topic on its own, we shall restrict ourselves to a sketch of the analogy between adaptive consequence when the underlying structured premise-set can be extended to an information-ordering, and Crispin Wright's proposal for the adoption of superassertibility as a truth-predicate (Wright, 2001). The first part of the analogy goes as follows. Wright claims that a minimal truth-predicate should both validate [DS] and [ASS].

[DS] 'p' is true iff p,

[ASS] to assert that p is to assert that p is true,

From the latter we derive as a feature of veridicality that "to be informed that p is to be informed that p is true". The second part of the analogy shows that if superassertibility is understood as:

> the property of being justified by some (in principle accessible) state of information and then remaining justified no matter how that state of information might be enlarged upon or improved. (Kvanvig, 1999)

We seem to have reasons to suppose that the interpretation of $i \Vdash [A]\phi$ as being informed that ϕ does tell us something about the veridicality of ϕ in those cases where $i \in \mathsf{INF\text{-}STAT}$. This, however, not in the sense that ϕ is veridical iff $i \Vdash [A]\phi$ (for $[A]\phi \vee [A]\neg\phi$ does not hold at all i for any ϕ), but in the weaker sense that if $i \Vdash [A]\phi$, then ϕ is true. Interestingly, the intrinsic properties of being informed seem to capture the greater part of the intuitions underlying [ASS], whereas [DS] finds a home in its extrinsic properties. The "anti-realist" interpretation of $[A]\phi$ can, roughly speaking, be situated somewhere between the concepts of assertibility and superassertibility. It certainly is — as it should be — a gappy property, and thereby shares a family resemblance with assertibility (on the gappiness of assertibility I side here with Wright and Restall (2001), and contra Tennant (1995) and Shapiro & Taschek (1996)). However, $[A]\phi$ also retains some of the stronger properties of superassertibility, most importantly the lack of reference to some maximal information-state, and the fact that a warrant for p is a warrant for the superassertibility of p only if p is superassertible, that is *stable* (i.e. if $i \Vdash p$, then $i \Vdash [A]p$ — a property satisfied by adaptive generated submodels if $i \sqsubseteq i^*$, that is if $i \sqsubseteq w$ and $w = w^*$, the formal version of the stability requirement).

Anyway, it is not too much to claim that to the anti-realist $[A]\phi$ is not silent about truth. There is no reason whatsoever to reject an interpretation of the presuppositions behind $[A]\phi$ (i.e. being undeniable) in terms of truth, viz. by holding that being informed that p is being informed of the veridicality of p. So, in a sense the anti-realist is free to speak of the intrinsic properties of veridicality. Contrary to the realist, she can coherently describe truth from an agent's perspective, but the question remains whether this move can be imposed upon her. Our answer to that question is unambiguous; she is free to go either way. If she refrains from using truth when applied to some limited agents, she obviously is — as an anti-realist — entitled to truth-free talk. If, on the other hand, she has to answer realistically inspired objections, she has all the tools to do so. This concludes the weak argument for a truth-free analysis of internal states.

5. *Concluding Remarks*

As this paper mixes the rather technical problem of applying an adaptive consequence-relation to structured premise-sets with the more philosophical issue of what it means to "hold the information that ...", or to "be informed that ...", both sides should be addressed and reconnected in this conclusion.

For the purely formal part of this paper, a brief comment suffices. The combination of generated submodels with the preferential models used in adaptive logics is basically a new one. More can certainly be obtained from the semantically oriented approach to adaptive modal logic initiated in this paper, and these mechanisms deserve a more in-depth study.

With respect to the application of the modal operator $[A]$ as a means to express the internal state of holding information, more should be said. A first remark expands on the distinction between "holding the information that p", and "holding p to be information" outlined in Floridi (2006). More precisely the fact that both are indistinguishable from an intrinsic point of view should be clarified. That is, $[A]p$ describes the internal state which is common to both notions. This, however, does not imply an endorsement of a rejection of the latter, stronger, concept. Throughout the analysis it was clearly shown that being in a veridical state is a property which cannot be verified nor falsified by an internal state (though, to the anti-realist the internal state is not necessarily indifferent towards veridicality). The conclusion should therefore not be the standard sceptic reaction, but only a clear understanding of the level at which full-blown veridicality obtains (or fails). In that sense, the model explains why the rejection of the possibility of being informed cannot be derived from the properties of the internal state (an often made mistake).

The second comment is more general and addresses the comparison between the adaptive approach, and the logic presented in Floridi (2006). Comparing them at the formal level is rather difficult as there is no obvious translation of expressions in one system into expressions of the other (this problem is addressed in Allo (in preparation)). On the interpretational level there is more to be said since basically both logical systems agree on the properties of "being informed". Concretely, every property of holding information in the sense of $i \Vdash [A]p$ and $i \in \mathsf{INF\text{-}STAT}$ is a property of holding information as $\Box p$ (where $\Box$ is a KTB-modality). The interesting point is that (i) not every property of $\Box p$ is a property of $i \Vdash [A]p$ and of $i \in \mathsf{INF\text{-}STAT}$; (ii) some properties of $\Box p$ are strictly a property of $i \Vdash [A]p$ (e.g. the lack of reflective states), while other properties of $\Box p$ are strictly a property of $i \in \mathsf{INF\text{-}STAT}$ (e.g. veridicality); and (iii) some properties of $\Box p$ are weakly endorsed by $i \Vdash [A]p$, but strongly endorsed by $\rho_g \subseteq \lambda$ and $i \sqsubseteq i^*$(e.g. constructibility, consistency, and on the anti-realist plan also stability).

ACKNOWLEDGEMENTS

Parts and early ideas which laid the grounds for this paper were presented at the VlaPoLo 9 Workshop (Ghent, 2004) as *Structured Proofs and Adaptive Logic*, and at the UNILOG-Conference (Montreux, 2005) as *Labelling in Modal Adaptive Logics: An Example*. Research for this paper was largely finalised during a stay at the IEG, Oxford University – UK (funded by the FWO-Vlaanderen), where it benefited from discussions with Luciano Floridi.

The author also wishes to thank Sonja Smets and Jean Paul Van Bendegem for comments on an earlier version of this paper.

Centre for Logic and Philosophy of Science
Free University of Brussels
E-mail: patrick.allo@vub.ac.be
http://homepages.vub.ac.be/~pallo/

REFERENCES

Allo, P., 2005, Being Informative – Information as Information Handling, in: *WM2005: Professional Knowledge Management Experiences and Visions*, Althoff, K.-D., A. Dengel, R. Bergmann, M. Nick and T. Roth-Berghofer, ed., DFKI Gmbh, Kaiserslautern: 579–586.

Allo, P., in press, Formalising Semantic Information. Lessons from Logical Pluralism, in: *Computing, Philosophy, and Cognitive Science*, G. Dodig-Crnkovic and S. Stuart, eds., Cambridge Scholars Press, Cambridge: xxx–xxx.

Allo, P., in preparation, A Non-Preferential Semantics for the Logic of "Being Informed".
Batens, D., 2000, Minimally Abnormal Models in Some Adaptive Logics, *Synthese*, 125(1): 5–18.
Batens, D., 2001, A General Characterisation of Adaptive Logics, *Logique & Analyse*, 44(173-175): 45–68.
Batens, D., in press, A Universal Approach to Adaptive Logics, in: *Directions in Universal Logic,* Béziau, J.-Y., ed., Birkhauser, Basel.
Blackburn, P., M. De Rijke and Y. Venema, 2001, *Modal Logic*, Cambridge University Press, Cambridge.
Dunn, J. M., 1993, Star and Perp: Two Treatments of Negation, *Philosophical Perspectives*, 7: 331–357.
Dretske, F., 1999, *Knowledge and The Flow of Information*, CSLI, Stanford.
Floridi, L., 2005, Is Information Meaningful Data?, *Philosophy and Phenomenological Research*, 70(2): 351–370.
Floridi, L., 2006, The Logic of Being Informed, *This Volume*.
Hand, M., 1999, Antirealism and Falsity, in: *What is Negation?*, D. M. Gabbay and H. Wansing, eds., Kluwer Academic Publishers, Dordrecht: 185–198.
Hughes, G. E., 1975, B(S4.3, S4) Unveiled, *Theoria*, 41: 85–88.
Kessler, G., 1975, A Note on Future Branching Time, *Theoria*, 41: 89–95.
Kvanvig, J. L., 1999, Truth and Superassertibility, *Philosophical Studies*, 93(1): 1–19.
Meheus, J., 2002, An Adaptive Logic for Pragmatic Truth, in: *Paraconsistency, The Logical Way to the Inconsistent*, Carnielli, W., M. Coniglio and I. D'Ottaviano, ed., Marcel Dekker, New York: 167–185.
Meheus, J., 2006, An Adaptive Logic Based on Jaśkowski's Approach to Paraconsistency, *Journal of Philosophical Logic*, 35(6): 539–567.
Meheus, J., in press, Discussive Adaptive Logics: Handling Internal and External Inconsistencies, *Poznan Studies*, to appear.
Restall, G., 1999, Negation in Relevant Logics, in: *What is Negation?*, D. M. Gabbay and H. Wansing, eds., Kluwer Academic Publishers, Dordrecht / Boston / London: 53–76.
Restall, G., 2001, Constructive Logic, Truth and Warranted Assertibility, *The Philosophical Quarterly*, 51(205): 474–483.
Shapiro, S. and W. W. Taschek, 1996, Intuitionism, Pluralism, and Cognitive Command, *Journal of Philosophy*, 93(02): 74–88.
Stalnaker, R., 2006, On Logics of Knowledge and Belief, *Philosophical Studies*, 128(1): 169–199.
Tennant, N., 1995, On Negation, Truth and Warranted Assertibility, *Analysis*, 55(2): 98–104.
Wright, C., 2001, Minimalism, Deflationism, Pragmatism, Pluralism, in: *The Nature of Truth*, M. P. Lynch, ed., MIT, Cambridge: 751–787.