

THE INFORMATION IN INTUITIONISTIC LOGIC

Johan van Benthem, Amsterdam & Stanford, <http://staff.science.uva.nl/~johan/>

May 2008, *SECOND DRAFT*

To appear in *Synthese* special issue on ‘Philosophy of Information’ edited by Luciano Floridi & Sebastian Sequoiah-Grayson (Oxford)

Abstract Issues about information spring up wherever one scratches the surface of logic. Here is a case that raises delicate issues of ‘factual’ versus ‘procedural’ information, or ‘statics’ versus ‘dynamics’. What does *intuitionistic logic*, perhaps the earliest source of informational and procedural thinking in contemporary logic, really tell us about information? How does its view relate to its ‘cousin’ epistemic logic? We discuss connections between intuitionistic models and recent protocol models for dynamic-epistemic logic, as well as more general issues that emerge.

1 Logic and information, a web of story lines

The notion of information is not a standard theme in logic, but it really lies everywhere close to the surface. Last year, editing the *Handbook of the Philosophy of Information* with Pieter Adriaans, it fell to me at a late stage to help write a chapter on this theme (van Benthem & Martinez 2007). I quickly found an ‘embarras de richesse’: many natural notions of information occur entangled in logic, and their relations are often unclear. Epistemic logic describes information of agents in terms of *ranges* of possible worlds, and its dynamic versions describe the informational processes that update these ranges: observation, conversation, and the like. Situation theory focuses on *correlations* between local states of the world, and how agents stay attuned to these. These two semantic views merge naturally in what might be called ‘temporal event models’, but they do not exhaust the scene. Logic also has a pervasive third, proof-theoretic or algorithmic notion of information as *code*, manipulated by dynamic processes of ‘elucidation’, such as proof or computation. These do not produce information in the same sense as semantic update, since they work on more finely structured data, often represented in formal or natural languages. The relation between semantic and proof-theoretic accounts of information is a major challenge, though there are promising attempts at the abstraction level of relevant and categorial logics. In the smaller compass of logic, this diversity reflects the many

complementary stances in science in general, from logic to Shannon information theory or Kolmogorov complexity, whose integration raises intriguing foundational questions for the philosophy of information as put on the map in the cited Handbook.

But the immediate impetus for this short note comes from Sebastian Sequoiah-Grayson, who observed the same diversity within my own mind! Over the years, I have written about information in different frameworks, without checking for consistency, let alone reconciliation. Is there any thread linking the categorial Lambek Calculus as information theory (van Benthem 1991), modal and temporal logics of information (van Benthem 1989, 1996B), and dynamic epistemic logic (van Benthem 2006B)? My original plan for this paper was to show there was, but the grand story is too large, and still in the clouds.

So, here is my actual theme, having more to do with Dutch patriotism. Toward the end of the above-mentioned Handbook chapter, there is a short passage on what may be the earliest logical system based on considerations of information, *viz. intuitionistic logic*. And before doing anything else, it seemed of interest to me to at least get clear on that. What vision of information underlies intuitionism, and what issues does this raise?

2 Intuitionistic logic and information

Constructive proof Intuitionism views mathematics as a web of constructive proof patterns and matching definitions of objects. Thus, it infuses the logical constants of propositional or predicate logic with a proof-theoretic spirit, summed up in the famous Brouwer-Heyting-Kolmogorov ‘provability interpretation’ (van Dalen & Troelstra 1988). For instance, a proof of a conjunction $A \wedge B$ is a pair of proofs for A and for B , and a proof for an implication $A \rightarrow B$ is a method for transforming proofs of A into proofs of B .

Semantics Over time, intuitionistic logic has picked up semantic models of independent interest, starting with algebraic and topological interpretations in the 1930s. Then, Beth 1956 proposed models on topological Baire space, or equivalently, over trees of finite or infinite sequences. These trees model abstract processes of investigation, and in line with the proof idea, intuitionistic formulas are true at a node when they are ‘verified’ there. By now, the dominant version of this idea are intuitionistic Kripke models (Kripke 1963), which we will take here as *partial orders* $\mathbf{M} = (W, \leq, V)$ with a valuation V , setting:

$$\begin{aligned}
\mathbf{M}, s \models p &\text{ iff } s \in V(p) \\
\mathbf{M}, s \models \varphi \wedge \psi &\text{ iff } \mathbf{M}, s \models \varphi \text{ and } \mathbf{M}, s \models \psi \\
\mathbf{M}, s \models \varphi \vee \psi &\text{ iff } \mathbf{M}, s \models \varphi \text{ or } \mathbf{M}, s \models \psi \\
\mathbf{M}, s \models \varphi \rightarrow \psi &\text{ iff for all } t \geq s, \text{ if } \mathbf{M}, t \models \varphi, \text{ then } \mathbf{M}, t \models \psi \\
\mathbf{M}, s \models \neg \varphi &\text{ iff for no } t \geq s, \mathbf{M}, t \models \varphi^1
\end{aligned}$$

Here, in line with the idea of ‘accumulating certainty’, the valuation is ‘*persistent*’:

if $\mathbf{M}, s \models p$, and $s \leq t$, then also $\mathbf{M}, t \models p$.

There is no revision *en route*. The truth definition lifts this behaviour to all complex formulas. E.g., a negation says the formula itself will never become true at any further stage of the process. This explains why Excluded Middle $p \vee \neg p$ is invalid, since this formula fails at states where p is not yet verified, though it will later become so. This may happen in situations of several types: compare the black dots in the two pictures below, which stand for the beginnings of informational process that unfold as downward trees.

We will discuss their differences in terms of dynamic information flow below:²



Here, one might criticize intuitionistic negation for failing to distinguish between two notions that are not obviously the same: (a) global *future absence of verification*, and (b) the intuitively more primitive notion of *local falsification*, which has been emphasized by Wansing, Gurevich and many others (cf. Wansing 1996). Indeed, at end-points in the above structures (leaves of the pictured trees), the above semantics makes a volte-face, reading lack of further verification as falsity in the current situation: something which seems contrived. I am in sympathy with this criticism, but will not pursue it here.

¹ The technical differences between Beth and Kripke models will not concern us here.

² For instance, unlike the first, the second process reaches an end-point where $\neg p$ holds.

Interpreting the models Now, what do these semantic models really stand for? It is of interest to read sources like van Dalen and Troelstra 1988 or Dummett 1977 to see what they say on this crucial matter. Indeed, they do not address this issue in detail at all. But they do say informally, by way of motivating the mathematical formalism, that a model describes a process of investigation by an agent, and that its worlds are ‘information situations’ or information stages’ where certain ‘knowledge’ has already been achieved.³ One reason why the story remains thin is that, in intuitionistic logic like in many other areas of logic, semantics followed proof theory – and models were devised to make sure that an already given language got a meaning, and some given proof calculus turned out complete. The very term *semantics* has this flavour: models come in the service of some given practice. But one can also take the opposite viewpoint, with semantic structures themselves as the primary object of study. That is what we will do in this paper.⁴

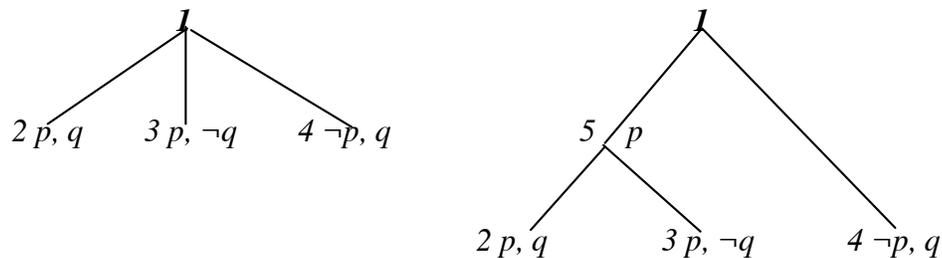
To return to our main concern, what notion of information is represented by intuitionistic models, of the Beth or Kripke variety? Henceforth, we just think of branching *tree-like*

³ Dummett even gives two different readings of the turn-style \models , which we forego here.

⁴ This may be the more common stance in science. If the above structures really model information processes, then we want to study them for their own sake – and only ask later what language can describe them, and what sort of logic then emerges. But even with models originally invented as semantics for some language, one must always ask if that language is really the most appropriate medium for describing them. There is a healthy spiral of iterative redesign for both languages and models, with two independent levels. Indeed, for intuitionistic logic, it has been asked why it sticks to the vocabulary of classical logic: tacitly, and perhaps infelicitously, adopting ‘the other’s turf’ for stating its own basic tenets and its deviant views. The same distinction plays in philosophical logic. Some people say we are just *analyzing meanings* of key notions in logic: natural language is given first, and it has a stable semantic content that just needs to be brought forth. I myself think such stable linguistic usage largely an illusion, and we are mostly engaged in *conceptual analysis* and Carnapian ‘rational reconstruction’ at the level of models, fitting our language later to our design for the task at hand. In particular, intuitionistic logic is not a theory of meaning of logical constants, alternative to their ‘classical meanings’. It is rather an analysis of information and constructive reasoning, and the intuitionistic language should be the best system of ‘logical switches’ that make such reasoning work – even if this means extending the classical languages of logic.

models, because they are closest to the intuitionistic tradition, and most congenial to our discussion. To make things easier, let us assume the trees are finite.⁵ Intuitively, each branching tree describes an *informational process* where an agent learns progressively about the state of the actual world, encoded in a propositional valuation. At end-points of the tree, all information is in, and the agent knows all the facts about the actual world.

Process equivalence and bisimulation A lot can be asked about such models without yet having any language at all, just given their process flavour. In particular, when are two informational processes of this sort equivalent? When the set of propositional valuations on the leaves is the same? Or when they are similar in some stricter sense, taking into account how the intermediate stages lie? As a concrete illustration, in the former sense, the following two information models are equivalent, while in the latter, they are not:



Intuitionistic formulas can easily differentiate between the roots of these two models,⁶ since they side with modal logic on the finer structural equivalence of *bisimulation*, which records choices at intermediate stages.⁷ Bisimulation is a standard equivalence from process theory, but there are others. Indeed, it seems a legitimate, though neglected, foundational issue deserving explicit discussion which notion of process equivalence is most appropriate to information gathering over time.

⁵ Everything I say makes sense for trees with *infinite branches*, too, but I want to keep things light. Even so, I may be skipping a finesse here. As dan Isaacson pointed out: infinite Beth models and finite Kripke models take a slightly different view of things that will be relevant below.

⁶ The simple verification, as well as later elementary observations of this sort are left to the reader.

⁷ For a definition, we refer to the standard literature: e.g., Blackburn, de Rijke & Venema 2000.

Indeed, all intuitionistic models are equivalent to tree models, up to bisimulation – though this format also suppresses some interesting structure: see our discussion of process types in Section 3.

A crucial distinction But the main point of our discussion here is rather the following fundamental distinction concerning information which comes to light here. Intuitionistic logic really registers two basic kinds of information:

- (a) ***factual information*** about how the world is; but on a par with this:
- (b) ***procedural information*** about our current investigative process.

To me, the latter feature concerning the way we learn the facts is the real semantic side of ‘constructivism’. How we get our ‘knowledge that’ matters deeply, and while the leaves record factual information, the branching structure of our tree models, and in particular, available and missing intermediate points, encodes agents’ knowledge of the latter kind.⁸
⁹ What is more, the distinction between factual and procedural information is not just an ad-hoc invention in our setting here: it makes sense everywhere, as soon as you grasp it.

3 Modal logic and information

Intuitionistic logic and ‘information loading’ We said before that providing semantic models for a language does not imply that the language is best for those models. Which language best describes the informational processes that intuitionism is, or should be, concerned with? Intuitionistic propositional logic focuses exclusively on stable knowledge achieved by a certain stage, no matter how the process continues from there on. And it does so in a rather peculiar way, viz. by ‘loading’ the account of the logical constants. Thus, ‘*not-A*’ does not only say that we do not have *A* right at the current stage,

⁸ Tim Williamson has asked whether the distinction ‘factual’/‘procedural’ is really robust as a feature of information. I tend to think it is context-dependent, but illuminating all the same.

⁹ As it stands, intuitionistic logics puts very few constraints on informational processes. For instance, different branches of the tree may have very diverse ‘leaps’ of informative value from node to node. Indeed, if we make the trees more uniform, say, by requiring that the admissible steps are precisely those corresponding to getting to know just one atomic proposition at a time, we get stronger logics, describing particular types of process. This perspective might give a better view of so-called ‘*intermediate logics*’ in between intuitionistic and classical logic. Usually just a belt of ‘asteroids’ in between two planets, these may now arise as well-motivated information theories.

but also that we will *never* get it, where ‘never’ refers to the future of the current process. Van Benthem 1993 calls this ‘information loading’ feature of intuitionism *implicit knowledge*, and contrasts it with the *explicit knowledge* found in *epistemic logic*, which keeps the classical account of the Boolean operations, but adds an explicit modal operator

$K\varphi$ saying that the agent knows that φ , or semantically, $\mathcal{M}, s \models K\varphi$
iff φ is true in all worlds t epistemically accessible from s in \mathcal{M} .

Of course, loading allows the intuitionistic language to express interesting notions. For instance, the formula $\neg\neg\varphi$ says that every stage has a later stage where φ is true, or in finite trees, each end node makes φ true. This is close to saying that, barring details of the process structure, we ‘know’ now already that φ must be factually true. Intuitionism distinguishes this from having verified φ per se, a distinction to which we will return.

Modal logic and ‘unloaded’ explicit information structure By their loading of the classical logical constants, intuitionistic formulas describe the structure of an information process in a somewhat roundabout way. A somewhat more straightforward alternative was proposed in van Benthem 1989, 1996, putting forward a simple view of *modal logic* as a theory of information. We let the universal modality say the following:

$\mathcal{M}, s \models []\varphi$ iff φ is true at all future states $t \geq s$,

and then express everything in the earlier truth definition by letting classical constants just stand for their old denotations, while reading an intuitionistic negation as the modal combination $[]\neg\varphi$ of a classical and a temporal operator, and an intuitionistic implication as the modalized implication $[](\varphi \rightarrow \psi)$.¹⁰ This setting is almost the standard semantics of the modal logic *S4*, except for the fact that atomic propositions are upward preserved.

¹¹

¹⁰ This link to intuitionistic logic is nothing but the famous *Gödel translation* into *S4*.

¹¹ Thus, our complete logic is *S4* plus additional atomic persistence laws $p \rightarrow []p$.

Incidentally, this does mean that the set of validities is not substitution-closed.

The modal language is more expressive, however, since it can also make further types of assertions about the current information stage. Typically, *existential modal formulas* $\langle \rangle \varphi$ say that the formula φ may still become true later, which may be useful as intermediate process assertions just like the “maybe”s that we employ in ordinary language. Unlike intuitionistic assertions, however, these are not persistent, as new facts moving us higher up in the tree may preclude φ from becoming true. Modulo a few technicalities, the intuitionistic language is the ‘persistent fragment’ of the modal one, and hence the two views of informational processes go well together as one *intuitionistic/modal paradigm*.

Modal logic as information theory Here are a few further features of the framework.

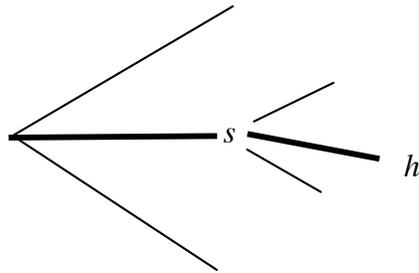
First, it is often said that the above means, without further ado, that intuitionistic logic ‘lies embedded in the epistemic logic *S4*’, which differs from the usual epistemic logics in having positive introspection (the axiom $[]\varphi \rightarrow [][]\varphi$), but not negative introspection, which would be the invalid *S5*-axiom that $\langle \rangle \varphi \rightarrow []\langle \rangle \varphi$. This is extremely misleading. The modal operator packages factual information about endpoints with something else: viz. the above procedural information about the current process leading up to there. In this light, the failure of negative introspection is wholly understandable as coming from the temporal logic, rather than from our bare access to the factual information content.

In general, the modal language affords a more explicit view of further tree structure that one might want to impose. Suppose we want the process to have a uniform description in terms of available information-producing moves. Say, we can make observations about proposition letters one by one, determining whether they are true or false. In that case, some intuitionistically acceptable trees are ruled out as possible process descriptions. For instance, we cannot have an initial node without p, q true go immediately to an end-point having both p, q true – and this missing model will be reflected in new validities beyond modal *S4*, say, $(\neg p \wedge \neg q \wedge \langle \rangle (p \wedge q)) \rightarrow \langle \rangle (p \wedge \neg q) \vee \langle \rangle (q \wedge \neg p)$. Indeed, it is easy to construct more sophisticated examples of constraints on trees arising in this way.¹²

¹² There are some complications here. The logic of the particular restriction on atomic observations is not substitution-closed, since it treats proposition letters differently from complex formulas. In terms of more general restrictions leading to intermediate logics – cf. Footnote 8 – the situation is

Finally, the modal view set forth here is related to earlier information-based models, such as the ‘data semantics’ of Veltman 1984, Landman 1986. But there is a difference. Data semantics is about universes of *information pieces* that stand to each other in an inclusion relation. However, there is no coherent view of the worlds s in the above models as ‘information pieces’.¹³ Our worlds are *stages* of an investigation, but these do not admit of a natural calculus of addition, and even the associated sets of formulas true at worlds, another candidate for ‘information pieces’, do not seem to have that structure either.¹⁴

Coda: decoupling factual and procedural information in temporal logics Our crucial distinction between procedural and factual information lives inside our modal language, since the $[]$ operator involves both. But this intriguing mixture may also be a source of confusions. At least, there is a wide-spread alternative for describing information trees, treating them in just the same way as any other process, viz. in an branching *modal-temporal* language. Here, the picture is again the familiar one of *branching trees*:



But now, we evaluate formulas at points s on histories h (complete branches through the tree), splitting our language more radically into modal and temporal aspects:

made somewhat uninteresting by our restriction to *trees*. E.g., the natural schematic constraint of ‘Weak Excluded Middle’ $\neg\varphi \vee \neg\neg\varphi$, which expresses ‘convergence’ over arbitrary Kripke models, would force our trees to be linear chains, making the information process somewhat trivial.

¹³ Analogies with categorical or relevant logics of information pieces remain strong, but just *how*? Likewise, connections with Scott information systems remain tantalizing.

¹⁴ Over arbitrary Kripke models, instead of just trees, moving to a view of states as information sets rather than temporal stages, van Benthem 1996B proposes two new binary modal operators for the natural lattice operators of *supremum* and *infimum*. These correspond roughly to taking deductive closures of unions and intersections, respectively, of factual theories associated with the nodes.

$$\begin{aligned} \mathbf{M}, h, s \models G\varphi & \quad \text{iff} \quad \text{for all } t \geq s, \mathbf{M}, h, t \models \varphi \text{ ('always in the future')} \\ \mathbf{M}, h, s \models []\varphi & \quad \text{iff} \quad \text{for all } k =_s h, \mathbf{M}, k, s \models \varphi \text{ ('necessarily now')} \end{aligned}$$

where $k =_s h$ if the two histories k, h are the same up to and including the point s . In this setting, the complete histories may be viewed as the possible outcomes of the process. The earlier modal box $[]\varphi$ says that φ is true at each stage on each future history: which would now be written explicitly as $[]G\varphi$, separating out modal and temporal aspects. We will return to this temporal perspective below, since it also lies beyond epistemic logic.

Remark on mathematical foundations We end with a digression. We have embedded intuitionistic logic in larger explicit modal-temporal theories of information processes. Does this make any sense from its original motivation, as an account of constructive mathematical reasoning? Or is it just a twist toward a latter-day theme: rational agency and information flow? I think it might be very interesting to rethink the issues here. Both classical and constructive meta-mathematics focus on mathematical theories, the *products* of some rational process of development that is left implicit. But if we were to study mathematical theories *together with* an explicit logical account of the dynamic processes producing them, the richer systems in this section might make a lot of foundational sense.

4 Epistemic logic and information

Now let us forget the intuitionistic perspective for a while, and briefly recall the basics of another paradigm for dealing with information flow, viz. *epistemic logic* in its modern dynamic sense (van Benthem & Martinez 2008, Baltag, van Ditmarsch & Moss 2008).

Base language and models The base language has operators $K_i\varphi$ for ‘agent i knows that φ ’, interpreted over models $\mathbf{M} = (W, \{R_i\}_{i \in I}, V)$ where the R_i are epistemic accessibility relations among the worlds for the agents.¹⁵ More precisely, as mentioned earlier:

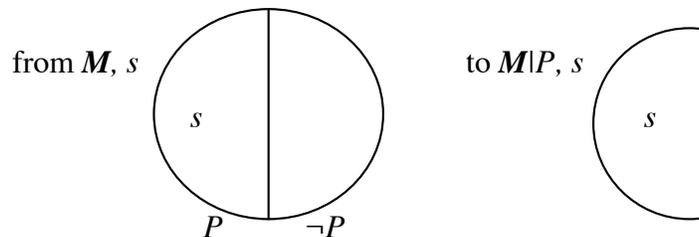
$$\mathbf{M}, s \models K_i\varphi \text{ iff } \mathbf{M}, t \models \varphi \text{ for all } t \text{ with } R_i s t.$$

¹⁵ ‘Worlds’ here can be as light as hands in your current card game, or the possible states of the light switches in your home. The old discussion about full-fledged metaphysical possible worlds which would not even make the most baroque episodes of Star Trek seems largely over, fortunately.

Details of this framework can be looked up in any standard text, and we will not dwell on them here, except to note that we will use equivalence relations, validating the system of ‘multi- $S5$ ’. Thus it seems that we are already far from the ‘epistemic logic’ behind intuitionism, which merely validated modal $S4$, endorsing just positive introspection. But this ‘ $S4$ versus $S5$ ’ is merely the usual ‘shallow analysis’, defining notions of knowledge or belief by the syntactic axioms they verify, rather than the deeper semantic conception underlying them. *Axioms are symptoms!* We will see later that ‘ $S4$ ’ and ‘ $S5$ ’ live in perfect harmony, once we ‘de-construct’ intuitionistic models in the right manner.

Another crucial observation right now is that we are not looking at epistemic logic as an account of the philosopher’s notion of *knowledge*. As argued in van Benthem 2006A, van Benthem & Martinez 2008, the operator K_i should really be read as “*to the best of agent i ’s information*”, viewing the accessibility relations R_i as defining agents’ current range of uncertainty, i.e., as information states in the folklore sense. This range view comes with another idea from the common sense, viz. that new information decreases the current range, while ideal information is just the singleton set $\{w\}$ with w the actual world. This brings us to the idea that this dynamics should itself be an explicit part of the logic.

Information dynamics: observation and communication For our purposes, it suffices to consider the so-called *logic of public announcements*: events $!P$ of new *hard information* which change irrevocably what I currently know. These can be linguistic communications from some perfectly reliable source, or public inter-subjective observations. Formally, such an event triggers a change in the current epistemic model (M, s) with actual world s . More specifically, $!P$ eliminates all worlds incompatible with P , thereby zooming in on the actual situation. Thus the current model (M, s) changes into its definable sub-model $(M|P, s)$, whose domain is the set $\{t \in M \mid M, t \models P\}$. In a picture, one goes



Typically, truth values of formulas change in such an update step: agents who did not know P now do after $!P$. This leads to subtle phenomena, but one can keep track of them in *public announcement logic PAL*, extending epistemic logic with action expressions:

Formulas	$P:$	$p \mid \neg\phi \mid \phi \vee \psi \mid K_i\phi \mid C_G\phi \mid [A]\phi$
Action expressions	$A:$	$!P$

The fundamental semantic clause for the dynamic action modality is as follows:

$$\mathbf{M}, s \models [!P] \phi \quad \text{iff} \quad \text{if } \mathbf{M}, s \models P, \text{ then } \mathbf{M} \setminus P, s \models \phi$$

As for the corresponding calculus of ‘hard information flow’, since these ideas are still less familiar than ‘hard core’ static epistemic logic, we note that *PAL* is axiomatized completely by the usual laws of epistemic logic plus the following *recursion axioms*:

$$\begin{aligned} [!P]q &\leftrightarrow P \rightarrow q && \text{for atomic facts } q \\ [!P]\neg\phi &\leftrightarrow P \rightarrow \neg[!P]\phi \\ [!P]\phi \wedge \psi &\leftrightarrow [!P]\phi \wedge [!P]\psi \\ [!P]K_i\phi &\leftrightarrow P \rightarrow K_i(P \rightarrow [!P]\phi) \end{aligned}$$

These axioms are the ‘recursion equations’ of public information flow, performing a step-by-step compositional analysis of epistemic effects of arbitrary events of incoming hard information. In particular, the final axiom relates the knowledge that agents get after receiving new information to conditional knowledge which they already had before.¹⁶ For details, cf. e.g., van Benthem 2006B, van Ditmarsch, van der Hoek & Kooi 2007.

Observational and inferential dynamics In concrete scenarios of information flow, observation is often deeply intertwined with *inference* (van Benthem 1996). But *PAL* only ‘dynamifies’ the former, not the latter events. Though the logic allows for inferences about effect of observations, it does not treat inference itself as a dynamic process. This

¹⁶ Strictly speaking, this validity assumes *perfect memory* and other idealized epistemic features of agents. Richer systems of *dynamic-epistemic logic (DEL)* exist in this spirit which can deal with information flow in more complex settings, such as card games, where not all players have equal observational access to the current event: say, your drawing a card from the current stack.

issue is related to the general difficulty of providing a satisfactory account of ‘inferential information’ (van Benthem & Martinez 2008, Sequoiah-Grayson 2008). There seems to be a challenge here, all the more so from the earlier intuitionistic perspective, which after all, started out with formal proof. A first solution in a *PAL* framework was proposed in van Benthem 2008B, inspired by ideas from Jago 2006, allowing for updates to be caused by observations, as well as by calls to inference rules. This issue will return below. But for now, we address another concern, closer to the semantic analysis of observations.

Saving the procedural information: dynamic-epistemic logic with protocols Though *PAL* looks both information-oriented and procedural, there is a flaw from the perspective of our earlier analysis of intuitionistic logic. Technically, working inside out, successive application of the reduction axioms will reduce any assertion in this dynamic-epistemic language to an equivalent static *purely epistemic formula*. Thus, only factual information remains, while the procedural information encoded in the dynamic modalities evaporates. While one can put a plausible positive spin on this reduction, we seem to be missing a semantic parameter: often called ‘protocol information’ about the temporal evolution of the relevant informational process. Not all observations might be available in any order, not everything that is currently true can be said in civilized conversation, and so on.

Accordingly, van Benthem, Gerbrandy, Hoshi & Pacuit 2008 propose a protocol-based version *TPAL* of public announcement logic describing temporal evolutions consisting of

a family of finite or infinite sequences of update events $!P$
over time, starting from some initial epistemic model.

Such structures are closely related to the *epistemic-temporal* models of Fagin, Halpern, Moses & Vardi 1995 and Parikh & Ramanujam 2002 (cf. van Benthem & Pacuit 2006), which have the same branching structure discussed at the end of Section 3 above.

We hope that this sketch suffices (the full machinery and results are in the cited paper). Now the earlier recursion axiom for knowledge after update changes to this new valid equivalence – where the crucial role of the procedural information now becomes explicit:

$$[!P]K_i\phi \quad \Leftrightarrow \quad \langle !P \rangle T \rightarrow K_i(\langle !P \rangle T \rightarrow [!P]\phi)$$

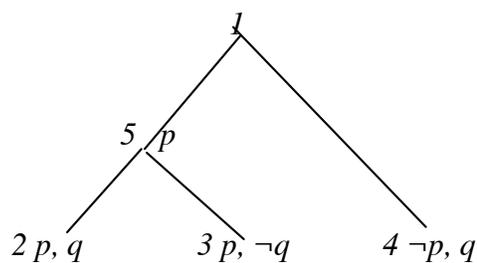
In words, the agent will know that ϕ after event $!P$ has taken place, iff, when the latter event is executable at all, he knows beforehand that, conditional on its executability, $!P$ will produce effect ϕ . This blocks the earlier epistemic reduction, since procedural assertions $\langle !P \rangle T$ about what can be observed or announced according to the current protocol will now remain crucial to ‘unpacking’ complex dynamic-epistemic formulas.

There are actually some variant $[!P]K_i\phi$ axioms for this logic, depending on how much agents are taken to know about the current protocol. In epistemic-temporal logics, that protocol is usually *common knowledge*, but in the most general version of *TPAL*, agents need not know the current protocol at all, which shows in the fact that different sequences of events $!P$ (including the empty sequence) may be attached to different accessible worlds in the initial epistemic model. These finer distinctions will return in what follows.

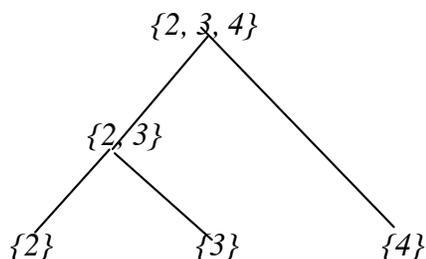
It may be clear where we are heading: dynamic-epistemic *TPAL* protocol models seems close in spirit to our earlier intuitionistic/modal models as investigative processes mixing factual and procedural information. This will be the subject of the following section.

5 From intuitionistic/modal to epistemic models of information

Introduction How would we model the intuitionistic/modal informational scenarios of Sections 2, 3 in dynamic-epistemic logic? Consider one of our trees again:



Epistemic logic assigns knowledge about some relevant situation in terms of possible worlds representing ways that situation might be. At stages of the tree, the obvious candidates for this are the end points below it (intuitionistically, these recorded complete valuations about the facts of the matter): or equivalently, the *complete histories* of the process. Thus, we can assign epistemic models at each node as follows:



One way of seeing this is as a family of epistemic models which decrease over time.¹⁷ Factual knowledge is then what is true in all current worlds. Moreover, the steps in the tree seem to correspond to epistemic events adding information about the current set of options. Let us now discuss this in more detail, fishing for more interesting things.

Trading future for current uncertainty in games The preceding is reminiscent of an issue in the epistemic analysis of *games*. In a game of ‘perfect information’, players know where they are at each node in the extensive game tree, but they do not know how the future will develop. But there is also a folklore observation that such ‘global’ uncertainty about the future can be converted into ‘local’ uncertainty about the present. The construction is simply this (cf. van Benthem 2004 for a slightly more formal version). Given any game tree G , assign epistemic models M_s to each node s whose domain is the set of complete histories passing through s (which all share the same past up to s), letting the agent be uncertain about all of them. ‘Worlds’ in these models may be seen as pairs (h, s) with h any history passing through s .¹⁸ It is natural to view the resulting structure $M(G)$ as a *TPAL* protocol model, where the actions are announcements which move is taking place. This will cut down the current set of histories in just the right manner.¹⁹

¹⁷ Actually, at some peculiar stages, no decrease may take place at all: see below.

¹⁸ We will ignore multi-agent issues, important though these are to games.

¹⁹ Actually, a game of perfect information becomes a so-called game of *imperfect information* in this manner, and this raises further issues. For instance, it is natural to distinguish two kinds of imperfect information in such games, *observation uncertainty* about how the game has developed so far, and *expectation uncertainty* about how it is going to continue. The former is the main sort of uncertainty analyzed in dynamic-epistemic and epistemic-temporal logics, and it is interesting to see what role it might play in an intuitionistic setting. Nodes in our trees might then have epistemic accessibility links, as in ‘intuitionistic modal logics’; cf. Bozic & Dosen 1984.

Languages This was about the structures, but what about matching *languages*? As usual with model constructions, the elegant language connection runs in the opposite direction. Consider a language over the induced models $\mathbf{M}(G)$ which has proposition letters p true at worlds (h, s) , Boolean operations, one-step modal operators $\langle !a \rangle \varphi$ for announcing that move a is played, plus an iterated version $\langle * \rangle \varphi$ of this saying that formula φ is true after some finite sequence of moves. Moreover, the language will have an epistemic operator $K\varphi$ quantifying over all currently accessible worlds as constructed above. The proper comparison for this formalism is clearly with a branching-temporal language associated with the game G , evaluating formulas φ as before in the format $G, h, s \models \varphi$.

Indeed, there is an obvious translation τ from the language of $\mathbf{M}(G)$ to that of G . Proposition letters remain the same, so do Booleans, announcement modalities $\langle !a \rangle$ go to action modalities $\langle a \rangle$, future modalities $\langle * \rangle$ go to the modal-temporal $\langle \rangle F$ (‘in the future on some currently indistinguishable history’) in the branching-time sense of Section 3, and finally, the epistemic modality K becomes exactly the local universal modality $[]$ of our earlier discussion of branching temporal models. In all, this translation τ supports the following equivalence, whose proof is a simple induction:

$$\mathbf{M}(G), (h, s) \models \varphi \text{ iff } G, h, s \models \tau(\varphi).$$

While this is pleasing, it may not satisfy readers expecting a translation in the opposite direction, reducing modal talk about games G to translated talk about the models $\mathbf{M}(G)$. But clearly, we can say at least this much: the central future modality $[]$ of the games, viewed as analogues of our intuitionistic-modal tree models, corresponds precisely with the operator combination $K[*]$ of the models $\mathbf{M}(G)$. A more precise discussion of languages appropriate to our game structures would also involve a deeper analysis of appropriate notions of *epistemic-temporal bisimulation*, which we forego here.²⁰

From intuitionistic-modal models to TPAL protocol models In principle, the same construction converts intuitionistic or modal tree models into *TPAL* protocol models. But there are also differences with the above scenario. First, we lack unique labels for the

‘moves’ in our tree: they are just anonymous upward inclusion links, which we can only identify by what is true at nodes after them. Also, we have no unique description of each history, since we need not (and in general, cannot) assume that different end-points in the tree carry different valuations.²¹ In other words, we want to ‘dynamify’ intuitionistic-modal models by making their hidden actions explicit, but what are they really?

We will discuss this in a moment, but for now, we note one further thing which makes the resulting models special among all *TPAL* protocol models. In general, a *TPAL* model may get stuck at the end of a history with some epistemic model whose worlds cannot be distinguished by further available events *!P*. But in our converted tree models, each individual world (read: history) is identified uniquely as the only ‘survivor’ at the end of its announcement history. We may call *TPAL* models like this *revealing*: a unique world always comes out. This reflects the initial motivation for intuitionistic models at the beginning of this paper: at end-points, they have identified the true world uniquely.

Announcement actions The first type of action naturally associated with transitions in intuitionistic-modal tree models are *public announcements*. Assuming that each end-point is uniquely definable in the modal language, each stepwise shrinking of the set of reachable endpoints can be defined as the announcement of the negations of all definitions for the endpoints that drop out. And in case there is no reduction in reachable endpoints (see below), we can still put the announcement *!T*. Disregarding the behaviour of proposition letters, this is probably all that can be said. Thus, to a first approximation,

*Intuitionistic logic describes effects of observations of facts,
but without making the nature of these observations explicit.*

²⁰ Extended bisimulations would also make sense for *TPAL* epistemic temporal protocol models.

²¹ Unique definability might have a bit more of a chance if we allowed past modalities in our epistemic temporal language, describing the unique histories that led up to the end-points.

There is more to be said here. Given the possible non-distinguishability of stages in the tree, available announcements may run short, and many technical issues can be raised.²² I leave these matters to another occasion, since they do not add to my present purposes.

Actions of explicit ‘seeing’ Much more intriguing, however, is the need for a second type of dynamic action beyond our *TPAL* models, which comes to light when we consider our initial examples in Section 2 of failure for the Double Negation law:



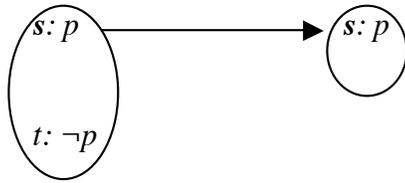
The second model M_2 poses no problems. We put different singletons at the end-points, and their union at the root. A protocol with just the announcements $!p$ and $!\neg p$ will ‘split’ these as required. But now consider the first model M_1 . One natural *TPAL* version would just put the same singleton set $\{s\}$ at both nodes, since nothing is ruled out going from one to the other. But then, no information can flow, and, if we think of p as a property of the end state, we already *knew that p* at the initial stage, by the semantics of the K operator. Knowledge of a formula in this sense, however, matches the intuitionistic operator $\neg\neg$, or its modal counterpart $[]\langle \rangle$, referring to eventual truth in all reachable end-points. But the point of intuitionistic models is that actually putting p at a stage means more than its ‘inevitability’ in this sense. What stronger event is taking place?

Before going there, let us avoid a confusion. While we can treat proposition letters true at end-points as true at the epistemic worlds in our *TPAL* model corresponding to that endpoint, proposition letters true at stages (not yet complete histories) had better be viewed as variants p^* , which imply, without being implied by epistemic formulas Kp .

Phantom worlds and partial protocols I can see two ways to go with this scenario. One adds new ‘phantom worlds’ to the epistemic models at each stage, which do not support

²² E.g., the construction may work better when we identify trees modulo bisimulation.

any events in the protocol (which has to be ‘local’ now in the earlier-mentioned sense). Thus, we might have the following structure:



Note that, in this way, the *TPAL* model is no longer revealing in the earlier sense. We no longer have Kp at the initial world s , even though the only available sequence of events will lead to knowledge that p . But actually, I incline towards the following alternative.

Explicit acts of seeing What intuitive event corresponds to *getting* a proposition letter p true at some stage in an intuitionistic model, even when we already had *knowledge of p* in the inevitability sense? To me, this is the event of seeing as in the exclamation “I see”: viz. a conscious realization of knowledge which I had only implicitly. But this calls for another modeling, more in line with the proposals in van Benthem 2008B:

We endow each world with a store of formulas already seen to be true at it, making new worlds consist of pairs (w, X) with X a set of formulas true at w . Now there are two kinds of ‘seeing’, one implicit, one explicit (cf. Barwise & Perry 1983 for discussion of the contrast). The implicit kind are our public announcements $!P$, which remove pairs (w, X) where w does not satisfy P in the current model. The explicit variant $+P$ makes the same removal, but it also adds the formula P to the X -component of all worlds that remain. Viewed in this sense, an atom p at some stage in an intuitionistic-modal model indicates that it is not only true, but also belongs to the explicit store. Of course, this would require extending *TPAL* to deal with explicit ‘storage worlds’, which has not happened yet.²³

I find this outcome truly amazing. In a plausible sense, intuitionistic logic gets connected to both dynamic-epistemic logic and more syntactic logics of explicit seeing and ‘awareness’, merging events of public observation with private acts of realization. I must leave further exploration of this duality to another occasion (cf. van Benthem 2008A).

In particular, the internal events described in this technical way may arise in any way that converts implicit knowledge into explicit knowledge: an explicit observation, a conscious inference, an act of introspection, and so on. Thus, they unify across a ‘genre’.

To conclude Many further issues will occur to the reader here, such as the appropriate choice of temporal languages, or the mechanics of converse transformations from *TPAL* protocol models into intuitionistic-modal ones. But I think that the links made here are enough to ponder already. In any case, my aim with this section has not been to reduce one paradigm to the other, but rather to see what happens when we put them side by side.

6 Conclusions, objections, and further issues

Information dynamics and rational agency This paper shows how to view intuitionistic logic as part of a larger epistemic-temporal logic of constrained informational processes. This observation is not new, and van Benthem 1989, 1996B already develop it in some detail. But we added that, in its insistence on both factual and procedural information, intuitionistic-modal logic fits particularly well with current protocol-based dynamic-epistemic logics. The latter systems, in my view, are stepping stones toward an account of rational agents that engage in actions of observation, inference, and communication on a par, and we have seen that intuitionistic-modal logic indeed goes beyond a dynamic logic of brute observation. A next step should be the treatment of inferential dynamics, a point to which we return in a moment. In addition, the agency view links intuitionism with theories of update, belief revision and learning.²⁴ What the original motivation for intuitionism might have to add to this theory of rational agency are processes less commonly studied, such as the dynamics of *definition*, and *creation of new objects*.

But how true is this to intuitionism? Our procedural analysis is really just one way of construing the dynamics in intuitionistic semantics, and Samson Abramsky and Tim

²³ Hoshi 2008 uses the syntactic power of local protocols in *TPAL* to model events beyond purely observational update, such as making inference steps.

²⁴ Van Benthem 1989 makes this point as well, though in a very different technical setting.

Williamson have pointed out others. For instance, another procedural reading for the above two-world model \mathcal{M}_1 is as follows:

One may reach the endpoint where p is true, but one may also get stuck forever at the first stage where p is not available.²⁵

As pointed out by Dan Isaacson, this alternative take becomes even more vivid in the *Beth model* associated with the Kripke model \mathcal{M}_1 , which would be an infinite p -less comb with p -teeth sticking out at each stage. While acknowledging all this, I would claim that (i) the procedural interpretation of intuitionistic models is not settled, and deserves more attention, (ii) this interpretation may allow for genuine alternatives, and (iii) the particular option I explore in this paper, with forced progression along nodes, is of interest in its own right, since it is highly appealing from a logical point of view to think of procedures that mix external acts of fact-finding with internal acts of inference, realization, or whatever ways our mind has of turning implicit into explicit knowledge.

Even so, what about proof? While the preceding seems largely true to me, the fact remains that intuitionistic logic arose as an analysis of mathematical proof. But our analysis in this paper has completely downplayed this aspect, going the opposite semantic way. There are many questions here to which I have no answer, merely some observations. First, there is still an intimate connection between *truth* in Beth or Kripke semantics for intuitionistic logic and *provability* in models used in completeness proofs. In particular, the modal accessibility relation in an intuitionistic Henkin model is just *set inclusion* between disjunction-splitting consistent sets. This supports the intuitive idea of information stages that really grow all the time²⁶ – though it has to said that the act of stepping from one deductively closed set to a proper extension in the Henkin model means acquiring an additional fact beyond drawing an intuitionistically valid conclusion.

²⁵ This seems closer to the dead-lock option in the ‘phantom world’ scenario mentioned earlier. More generally, the construal with an option of staying at some intermediate node forever is one of the motivations behind *domain theory*: see Abramsky 2008.

²⁶ By contrast, non-persistent existential modalities $\langle \rangle \varphi$ mess up this idyllic inclusion picture, while also leading to more complex interpretations in terms of non-provability.

Also, in our models, I do find it difficult to make immediate sense of the original idea of intuitionism as proof. Consider a Modus Ponens inference from φ and $\varphi \rightarrow \psi$ to ψ . The semantics has no dynamic counterpart to this, since it says that if both premises are true at a stage of the information process, then so is the conclusion.²⁷ On the other hand, the proof-theoretic interpretation naturally says that providing a proof x for φ and y for $\varphi \rightarrow \psi$ leads to the new action of providing a proof ‘ $x+y$ ’ for ψ . While this natural inferential dynamics has semantic counterparts in models for relevant or categorical logic (cf. Dunn 2008), these do not show up in our dynamic-epistemic analysis. Thus, a dynamic logic of inference steps, as hinted at above, seems a plausible addition.^{28 29}

Intuitionistic logic and multi-agent interaction But precisely through the notion of proof, intuitionistic logic seems to have even more radical connections with the theory of rational agency. Constructive proofs were analyzed as winning strategies in two-player argumentative dialogues as early as Lorenzen 1955, and this theme has turned out fruitful generally in the ‘game semantics’ of linear logic and programming (cf. Abramsky 2008). Moreover, the many links with games in the literature on dynamic-epistemic logic suggest that this link may be more than a coincidence, also in our semantic setting.

Frankly, I am somewhat undecided on what to make of this right now. It might be that intuitionistic logic points the way towards a grand synthesis of information analysis in the standard model-theoretic style with the dynamic view of logic as embodied in proof and games. Or it might also be that this meeting is a ‘lucky fluke’, since the language is so weak in expressive power that it happens to combine these different interpretations.

²⁷ Perhaps one might view an intuitionistic implication or a modalized $[[p \rightarrow q]$ as expressing a dynamic *dependency* between the two variables p, q : if one becomes manifest, so does the other.

²⁸ Cf. Jago 2007, Velazquez 2008, Hoshi 2008 for current attempts in dynamic logic style.

Three conclusions

This paper suggests three conclusions. First of all, intuitionistic logic remains a fascinating ‘implicit’ account of information, both factual and procedural, which suggests intriguing new connections to epistemic logic in its current dynamic guises as an ‘explicit’ account of information that agents may have and trade. Next, such a new take on intuitionistic logic suggests that the good old foundations of mathematics might at the same time be the most avant-garde theory of agency. And finally, and more generally, the case study of this paper confirms a pleasant suspicion: start with rethinking any major topic anywhere in logic, and you will find a theory of information underneath.

Acknowledgment I would like to thank Samson Abramsky, Georg Gottlob, Dan Isaacson, Sebastian Sequoiah-Grayson, and Tim Williamson for helpful comments.

References

- S. Abramsky, 2008, ‘Information, Processes and Games’, to appear in P. Adriaans & J. van Benthem, eds., *Handbook of the Philosophy of Information*, Elsevier Science Publishers, Amsterdam.
- A. Baltag, H. van Ditmarsch & L. Moss, 2008, ‘Epistemic Logic and Information Update’, to appear in P. Adriaans & J. van Benthem, eds., *Handbook of the Philosophy of Information*, Elsevier Science Publishers, Amsterdam.
- J. Barwise & J. Perry, 1983, *Situations and Attitudes*, The MIT Press, Cambridge (Mass.).
- J. van Benthem, 1989, ‘Semantic Parallels in Natural Language and Computation’, In H-D Ebbinghaus et al., eds., *Logic Colloquium, Granada 1987*, North-Holland, Amsterdam, p. 331 – 375.
- J. van Benthem, 1991, *Language in Action: Categories, Lambdas, and Dynamic Logic*, North-Holland Amsterdam & MIT Press, Cambridge (Mass.).
- J. van Benthem, 1993, ‘Reflections on Epistemic Logic’, *Logique et Analyse* 34, 5 – 14.

²⁹ I forego a comparison with the proof/category-theoretic perspective of Abramsky 2008, which rather identifies the ‘intrinsic dynamics’ of intuitionistic formulas with their role as objects in categories where proofs are morphisms.

- J. van Benthem, 1996A, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- J. van Benthem, 1996B, 'Modal Logic as a Theory of Information', in J. Copeland, ed., *Logic and Reality. Essays on the Legacy of Arthur Prior*, Clarendon Press, Oxford, 135 – 168.
- J. van Benthem, 2004. 'Update and Revision in Games', Lecture Notes, Department of Philosophy, Stanford University.
- J. van Benthem, 2006A, 'Epistemic Logic and Epistemology: the state of their affairs', *Philosophical Studies* 128, 49– 76.
- J. van Benthem, 2006B, 'One is a Lonely Number: on the logic of communication', in Z. Chatzidakis, P. Koepke & W. Pohlers, eds., *Logic Colloquium '02*, ASL & A.K. Peters, Wellesley MA, 96 – 129.
- J. van Benthem, 2007, Interview, in L. Floridi, ed., *Five Questions in the Philosophy of Computing and Information*, Automatic Press, Roskilde.
- J. van Benthem, 2008A, 'Merging Observational and Access Dynamics in Logic', *Chinese Journal of Logic*, Vol. 1:1, Institute of Logic and Cognition, Soon Yat-Sen University, Guangzhou.
- J. van Benthem, 2008B, 'Tell it Like It Is', *Journal of Peking University, Humanities and Social Science Edition*, No. 1, 2008, pp. 80 – 90.
- J. van Benthem, J. Gerbrandy & E. Pacuit, 2007, 'Merging Frameworks for Interaction: DEL and ETL'. *Proceedings TARK 2007*, University of Namur. Extended 2008 version with Tomohiro Hoshi, ILLC Amsterdam & Stanford University.
- J. van Benthem & M. Martinez, 2007, 'The Stories of Logic and Information', to appear in P. Adriaans & J. van Benthem, eds., *Handbook of the Philosophy of Information*, Elsevier Science Publishers, Amsterdam.
- J. van Benthem & E. Pacuit, 2006, 'The Tree of Knowledge in Action', *Proceedings Advances in Modal Logic*, ANU Melbourne.
- E. W. Beth, 1956, 'Semantic Construction of Intuitionistic Logic', *Mededelingen Koninklijke Nederlandse Akademie van Wetenschappen*, New Series 19:11, Amsterdam, 357 – 388.
- P. Blackburn, M. de Rijke & Y. Venema, 2000, *Modal Logic*, Cambridge University Press, Cambridge.

- M. Bozic & K. Dosen, 1984, 'Models for Normal Intuitionistic Modal Logics', *Studia Logica*, 4, 217 – 245.
- D. van Dalen & A. Troelstra, 1988, *Constructivism in Mathematics*, two volumes, North-Holland, Amsterdam.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, 2007, *Dynamic Epistemic Logic*, Springer, Dordrecht.
- M. Dummett, 1977, *Elements of Intuitionism*, Oxford University Press, Oxford.
- R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, The MIT Press, Cambridge (Mass.).
- T. Hoshi, 2008, 'Notes on TPAL (Public Announcement Logic with protocols) and Explicit Knowledge', Department of Philosophy, Stanford University.
- M. Jago, 2006, *Logics for Resource-Bounded Agents*, Dissertation, Department of Philosophy, University of Nottingham.
- M. Jago, 2007, 'Logical Information is Vague', Department of Philosophy, University of Nottingham, to appear in *Knowledge, Rationality and Action*.
- S. Kripke, 1963, 'Semantical Considerations on Modal and Intuitionistic Logic', *Acta Philosophica Fennica* 16, 83 – 94.
- F. Landman, 1986, *Towards a Theory of Information, the Status of Partial Objects in Semantics*, Foris, Dordrecht.
- P. Lorenzen, 1955, *Einführung in die Operative Logik und Mathematik*, Springer, Berlin.
- R. Parikh & R. Ramanujam, 2003, 'A Knowledge Based Semantics of Messages', *Journal of Logic, Language and Information* 12, 453 – 467.
- S. Sequoiah-Grayson, 2007, 'Information Gain from Inference', to appear in *Proceedings of LogKCA-07*, International Workshop on Logic and Philosophy of Knowledge, Rationality, and Action: San Sebastian, Spain, The University of the Basque Countries Press, Donostia.
- F. Velazquez, 2008, 'Inference and Update', ILLC, University of Amsterdam, to be presented at Workshop on Logic and Intelligent Interaction, ESSLLI Summer School, Hamburg.

F. Veltman, 1984, 'Data Semantics', in J. Groenendijk, Th. Janssen & M. Stokhof, eds., *Truth, Interpretation and Information*, Foris, Dordrecht, 43 – 63.

H. Wansing, ed., 1996, *Negation. A Notion in Focus*, Perspectives in Analytical Philosophy Vol. 7, de Gruyter, Berlin.