

LOGICAL PLURALISM MEETS LOGICAL DYNAMICS?

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

May 2008

Introduction

Where is logic heading today? There is a general feeling that the discipline is broadening its scope and agenda beyond classical foundational issues, and maybe even a concern that, like Stephen Leacock's famous horseman, it is 'riding off madly in all directions'. So, what is the resultant vector? There seem to be two broad answers in circulation today. One is *logical pluralism*, locating the new scope of logic in charting a wide variety of *reasoning styles*, often marked by non-classical structural rules of inference. This is the new program that I subscribed to in my work on sub-structural logics around 1990, and it is a powerful movement today.¹ But gradually, I have changed my mind about the crux of what logic should become. I would now say that the main issue is not variety of reasoning styles and notions of consequence, but the variety of *informational tasks* performed by intelligent interacting agents, of which inference is only one among many, involving observation, memory, questions and answers, dialogue, or general communication. And logical systems should deal with a wide variety of these, making information-carrying events first-class citizens in their set-up. This program of *logical dynamics* was proposed in van Benthem 1996. The purpose of this brief paper is to contrast and compare the two approaches, drawing freely on some insights from earlier published papers. In particular, I will argue that logical dynamics sets itself the more ambitious diagnostic goal of explaining *why* sub-structural phenomena occur, by 'deconstructing' them into classical logic plus an explicit account of the relevant informational events. I see this as a still more challenging departure from traditional logic. Diehard mathematicians still feel at ease with logical pluralism since it is all still a 'science of formal systems' describing 'inference', while to me, inference is just one way of producing information, at best on a par, even for logic itself, with others.

¹ Beale & Restall 2006 is a well-argued reference for the program of Logical Pluralism, including its 'parametrization' of logical systems in their choice of the relevant set of 'cases'.

1 Styles of reasoning

Structural rules Classical consequence $P \Rightarrow C$ from a finite sequence of premises P to a conclusion C says that C is true in every situation where all of the propositions in P are true. Without displaying any logical operations in the language at all, whether Booleans, quantifiers, or modalities, this bare-bones relation between abstract premises and abstract conclusions satisfies a number of interesting principles highlighted in Scott 1971 and later publications. These are the following *structural rules*:

<i>if $P, Q, R, S \Rightarrow C$, then $P, R, Q, S \Rightarrow C$</i>	<i>Permutation</i>
<i>if $P, Q, Q \Rightarrow C$, then $P, Q \Rightarrow C$</i>	<i>Contraction</i>
<i>$C \Rightarrow C$</i>	<i>Reflexivity</i>
<i>if $P \Rightarrow Q$ and $P, Q \Rightarrow C$, then $P \Rightarrow C$</i>	<i>'Cut'</i>
<i>if $P \Rightarrow C$, then $P, Q \Rightarrow C$</i>	<i>Monotonicity</i>

Together, these laws encode the basic 'style of reasoning' behind classical consequence. It treats the data that feed into a conclusion as sets (order and multiplicity do not matter), the inferential relation is a pre-order allowing for chaining of conclusions, and 'overkill' does not matter: accumulating more data is not going to endanger earlier conclusions.

The 'Bolzano Program' The 1970s and 80s saw a wave of other notions of consequence, reflecting quite different reasoning styles. Well-known are relevant logic (dropping monotonicity), default logics (dropping monotonicity and transitivity), resource logics in categorical grammar and linear logic (dropping contraction), and many others. Moreover, a general structure theory of these inferential relations developed in the work of Gabbay 1996, Dunn 1991, Restall 2000, and others, while Dosen & Schroeder-Heister 1993 coined the term 'sub-structural logics'. Van Benthem 1989 noted the analogy between this abstract level and the agenda for logic in Bernard Bolzano's *Wissenschaftslehre* (1837) which did not focus on 'logical constants', but on charting the key formal properties of different reasoning styles: deductive or probabilistic, in the common sense or according to strict philosophical standards. The term I proposed back then for this surprisingly modern enterprise: '*Bolzano's Program*', has never caught on, even though this original German-Italian pioneer continues to exert an appeal to logicians (van Benthem 1985, 2003A).

Representation theorems An analysis of logical systems in terms of mere structural rules has an attractive Spartan austerity, and the resulting hard-core principles also high-light surprising analogies across fields. Still, one wants to be reassured that this abstract proof theory has ties to some richer semantic or practical picture. This is what is provided by the usual *representation theorems*, of which there exists a great abundance – and I have dabbled with gusto in this cottage industry myself (van Benthem 1991, 1996A). Here is a ubiquitous folklore example, showing the semantic bare bones of classical consequence, writing bold-face \mathbf{P} for finite sequences of formulas as before:

Theorem An inference relation $\mathbf{P} \Rightarrow C$ satisfies the above five structural rules iff it can be represented by a map sending propositions P to sets $Set(P)$ with

$$P_1, \dots, P_n \Rightarrow C \text{ iff } \bigcap_{1 \leq i \leq n} Set(P_i) \subseteq Set(C).$$

Proof The proof is simply by setting $Set(B) =_{def} \{A \mid A \Rightarrow B \text{ in the given relation}\}$, and then checking that the given equivalence holds by an appeal to all given structural rules. ♣

More sophisticated representation theorems tie further notions of consequence to more elaborate semantic settings, beyond simple ‘set-intersection plus inclusion’ patterns. Some examples will be stated later in this paper. I will take it for granted that this wealth of notions and results provides a serious underpinning for logical pluralism today.

Two worries Even so, in Benthem 1989, I voiced two concerns about an exclusive focus on a level of abstract consequence. First, it seemed to be that many observations in terms of structural rules address *mere symptoms* of some more basic underlying phenomenon. For instance, non-monotonicity is like ‘fever’: it does not tell you which disease causes it. Thus, I was missing a deeper analysis of the underlying phenomena as a matter of logic.

Matching this was a second worry. Sub-structural logics often arise from ‘giving up’ some properties of classical consequence, while retaining the old formal language. But why not be radical *with respect to the language* as well, and reconsider what we want to say? Admittedly, this happened with linear logic and its ‘splitting’ of classical connectives, and the same is true to some extent for relevant logic as well. But, for instance, it has not happened with circumscription and default logics, and we will return to that issue below.

So, given this picture, and these concerns, can broad-minded logicians ‘dig deeper’?

2 Logical dynamics, rational agency, and intelligent interaction

Entanglement of informational tasks When asked to explain what logic is to a general audience, I often use the following scenario. You are in a café with two friends, where you have ordered a beer, a wine, and a water. Now some new person comes back with three glasses. What will happen? Everyone agrees that three things occur in sequence:

First the waiter asks “Who has the wine?”, say, and puts that glass. Then, he asks who has the beer, and puts that glass. And then, he does not ask any more, but just puts the remaining glass. Two questions, and then one inference!

When he puts that third glass without asking, you observe a logical inference in action: the information in the two answers received allows the waiter to just deduce where the third one must go. One can spell out this final stage in terms of a valid propositional schema

$$A \vee B \vee C, \neg A, \neg B \Rightarrow C,$$

whose power can be seen at work wherever people are solving *Sudoku* puzzles.

But to me, there is a unity to this scenario which gets torn when we just emphasize the final *inference*. The waiter first obtains the relevant information by *communication* and perhaps *observation*, and then, once enough data have accumulated, he infers an explicit solution. Now on the traditional line, only the latter deductive step is the proper domain of logic, while the former steps are at best ‘pragmatics’. But in my view, all these informational processes are on a par, and all should be within the compass of logic, which is about information flow in general, not just deductive elucidation. In my book, asking a question and understanding an answer is just as ‘logical’ an activity as drawing an inference. Thus, logical systems should account for both, as observation, communication, and inference occur entangled in most meaningful activities.² But what is involved in this program?

² This entanglement has a historical pedigree. Traditional Indian logic distinguished three principled ways of getting information. The easiest route is to observe, when that is possible. The next method is inference, in case observation is impossible or dangerous, as with a coiled object in a room where

Logics of inference and observation Our first task is to design richer logics where observations and inferences live on a par. For instance, the answer to the waiter’s first question reduces a space of 6 options to one with only 2, while the second answer reduces this to just 1, the correct assignment of drinks. Systems with explicit such steps exist by now in the form of *dynamic-epistemic logics*, which describe both the information which agents have at any given stage, and how their knowledge changes when they update their current state with new ‘hard information’ (say, an authoritative answer to some question). Also, complete axiomatizations exist for many such systems. Some details are provided below. These logics can be developed in a perfectly standard manner, so we really show that ‘logic can be more than it is’: the semantics and proof-theoretic techniques that we already possess can describe observation and communication as well as inference.³

Rational agents: from information to correction and purpose But even this is just a start. To me, modern logic is about *rational agency*.⁴ This view seems a natural culmination of a century of philosophical and computational logic after the great foundational era of the 1930s which focused on formal proofs without any agent at all – or at best, a computer as a single-agent device churning out new code. Rational agents have many ways of getting new information, and use it for a variety of purposes. Both inferential and observational skills are essential here, including observations of a non-public character, which permeate general communication, or playing games. But much more is involved! In particular, as has been realized since the 1980s, in addition to observation and knowledge update, there is the

we cannot see whether it is a piece of rope, or a cobra. And if these two methods fail, we can still resort to communication, and ask some expert. Similar ideas occur in medieval Western logic.

³ Admittedly, current dynamic-epistemic logics do not capture the dynamics of the waiter’s inference: what changes does it bring about [certainly, it does not change the final semantic state], and what are these good for? This famous difficulty of explaining ‘inferential’ versus ‘semantic’ information has no final solution that I am aware of. Van Benthem & Martinez 2007, van Benthem 2008A have more fine-grained combined accounts of observational and inferential dynamics.

⁴ This is not quite sharp, since I view dealing with irrational agents, and ‘bounded rationality’ as part of the logical story, too. Maybe ‘successful agency’ would be a better, more neutral term.

crucial formation of beliefs that guide our actions, and the associated processes of *self-correction* that revise beliefs when triggered by new information, ‘hard’ or ‘soft’.⁵

And this triad of observation, inference, and revision is still only part of the story. We can only make sense of actions by taking *preferences* into account, i.e., how agents *evaluate* situations. This is no ‘economic fad’: it is hard to think of meaningful communication or inference without keeping track of the ‘why question’ what it is good for.⁶

Intelligent interaction We are almost there, but not yet! Consider the waiter once more. Questions involve more than one agent, and their *social dynamics* involves higher-order knowledge which is crucial to communication. Asking you a normal question conveys that I do not know the answer, while I think you may. And your answer does not just transmit a fact, but it makes sure you know that I know, I know that you know that I know, and in the limit, it achieves *common knowledge*, a central notion in philosophy, linguistics, computer science, and cognitive science. It also underlies game theory, the best available current model of intelligent interaction in general. Indeed, the ability to move through an informational space keeping track of what other participants do and do not know, including the crucial ability to switch and view things from other people’s perspective, seems characteristic of human intelligence. Logic is just as much in the others as in ourselves.

This social interactive view goes back to the very roots of logic. While many people see Euclid’s *Elements* as the source, with its crystalline formal proofs and eternal insights, the true origin of the discipline may be closer to Plato’s *Dialogues*, an argumentative practice with clear patterns of confirmation and refutation between participants. It has been claimed that logic arose originally out of political and legal debate in all its three main traditions: Chinese, Indian, and Western. And this multi-agent interactive view has emerged anew in modern times. A striking example are the *dialogue games* of Lorenzen 1955, which recast

⁵ Revision is, of course, also the heart of the matter in more general processes of *learning*.

⁶ Even in mathematical proof by professionals, keeping track of *the point* is a crucial skill.

the foundations of logical consequence, explaining validity in terms of *winning strategies* for a proponent arguing the conclusion against an opponent granting the premises.⁷

The ‘Logical Dynamics’ program This grand program may be described more technically in terms of a distinction between *activities* and their *products*, which reflects a pervasive duality in natural language (van Benthem 1996A). ‘Dance’ is a verb denoting an activity, which then produces ‘dances’ like a tango or a waltz as a result. ‘Argument’ or ‘proof’ is a logical activity one can engage in, but it uses ‘arguments’ and produces ‘proofs’. Now traditional logic has emphasized products of logical activities, such as reasoning or seeing, while it has usually kept those activities themselves behind the scenes, as the motivating background stories. By contrast, the turn toward ‘logical dynamics’ is the conscious effort to make these activities themselves first-class citizens of logical theory. The systems described above are of this kind, but so are belief revision theory, ‘dynamic semantics’ (van Benthem, Muskens & Visser 1997), and many other research programs.

One way, not the only one, but certainly a good starting point, of creating such systems is by taking a close look at any existing system of logic, and asking: what are the dynamic activities or processes in the background? Thus, dynamic-epistemic logics arise from asking what processes would actually produce the models that standard ‘static’ epistemic logic is concerned with.⁸ By now, there is a wide range of ‘*dynamified*’ classical systems, again with the above-mentioned ones as examples. One virtue of this conservative approach is how it makes it clear from the start that the aim is not to do away with classical logics, but rather to enrich them and extend their natural boundaries.

Two worries Even so, many colleagues have a defensive response to this perspective, trying to fight some last stand of what defines ‘a logic’ as opposed to a dangerous outer

⁷ In my paper van Benthem 2008B, I sketch how current dynamic logics of information update, belief revision, preference change, and strategic interaction all address different aspects of this total behaviour, which must be integrated. ‘Logic, Rational Agency, and Intelligent Interaction’ (van Benthem 2007B) then lays out this modern version of Logical Dynamics in greater detail.

⁸ As a beneficial side-effect, such a model-transforming view also gives us a theory of systematic model construction for given epistemic scenarios, a topic usually left to ‘art’ and improvisation.

world. For instance (objections have varied over time), it has been said that ‘real logic’ must be about consequence relations, from which it follows that the dynamic perspective cannot be logic by definition, as it also includes information flow driven by observation and communication which is not naturally cast as ‘reasoning’. Others say that logic must be ‘the science of formal systems’ (a weird view, almost free of any exciting content whatsoever), making the ‘product’ view sacrosanct. While this is largely an issue of terminology, I find the objections interesting, because they themselves define logic in the non-dynamic manner which I find so limited. As a logician (I hope), I would say that the discipline of ‘logic’ is best described, not by any subject matter plus border patrols, but as that activity which is successfully performed by logicians using logical notions and tools, *wherever* those take them. Moreover, given the shifts in the historical agenda of the field, some modesty in claiming what logic *is* in some essentialist sense might be appropriate. But I hasten to add that not just logicians have this response. At a recent meeting with cognitive psychologists, I ran into opposition from the other side. When explaining that I wanted to investigate the logic of children’s mutual knowledge and interactive strategies in card games and other activities, plus the steps by which they make these more sophisticated in early childhood, I was told by eminent experts that this was not about reasoning couched *in language*, and so, it had nothing to do with logic, and I should leave this field to others.⁹ I mention these objections, not because they sway me, but to show my awareness that Logical Dynamics is one or maybe many bridges too far for most colleagues in the field.

The main theme of this paper Even so, I think that Logical Dynamics is the liveliest current alternative to Logical Pluralism. It, too, offers a coherent vision of what logic could become, and hence the real topic of this paper is to contrast and compare the two programs. We will do so in three case studies, starting from our earlier logics of observation.

3 Information update and consequence relations

As we have seen, even in basic scenarios of agency, inference and information update are intertwined. To get at that, in the above spirit of ‘dynamification’, we first need a good account of the ‘statics’ here. For that, we take standard *epistemic logic* of knowledge and

⁹ Just in case you are interested, the project *will* happen, and logicians *will* be involved.

related attitudes, including its current dynamic sense (Baltag, van Ditmarsch & Moss 2008). Following that, we move from the dynamic perspective to a consequence-based one, identifying a natural notion of ‘dynamic consequence’ which we then analyze in detail as a sub-structural logic in the ‘pluralist’ vein. Finally, we show how the two approaches are related, and how one might view the consequence approach as a way of seeking *natural abstraction levels* behind a given dynamic logic of some specific agent activity.

3.1 From static to dynamic epistemic logic

Epistemic base language and range models The propositional base language has operators $K_i\varphi$ for ‘agent i knows that φ ’, interpreted over models $\mathbf{M} = (W, \{R_i\}_{i \in \mathcal{A}}, V)$ where the R_i are epistemic accessibility relations among the worlds for the agents.¹⁰ More precisely:


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

In what follows, we will write $\langle K \rangle\varphi$ for the existential dual of this notion. Details of this framework can be looked up in any standard text, and we only note here that we will use equivalence relations for convenience, validating the logic of ‘multi-S5’. Another point is that we are not using epistemic logic as an account of the philosopher’s notion of *knowledge*. As argued in van Benthem 2006A, the operator K_i should 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., information states in the folklore sense. These ranges come with another common sense idea, viz. that new information decreases the current range, while ideal information is just the singleton set $\{w\}$ with w the actual world.

Information dynamics: observation and communication For our purposes, it suffices to consider the *logic of public announcements*: events $!P$ of new *hard information* which may change irrevocably what I currently know. These events 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 (\mathbf{M}, s) with actual world s . More specifically, $!P$ eliminates all worlds in \mathbf{M} that are incompatible with P , thereby zooming in

¹⁰ ‘Worlds’ here can be as light as hands in a card game, or the possible states of a traffic light.

on the actual situation. Thus the current model (\mathbf{M}, s) changes into its definable sub-model $(\mathbf{M}|P, s)$, whose domain is the set $\{t \in \mathbf{M} \mid \mathbf{M}, t \models P\}$. In a picture, one goes



Typically, truth values of epistemic formulas may change in such an update step: agents who did not know P now do after the event $!P$. This switching leads to subtle phenomena, but one can keep track of them in the following formalism.

Definition The language of *public announcement logic* *PAL* extends epistemic logic with action expressions denoting the preceding update steps: ¹¹

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:

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

When used with an epistemic logic for several agents, this language can also describe the effects of multi-agent conversation and communication. Indeed, it suffices for solving well-known puzzles like the ‘Muddy Children’ (Fagin, Halpern, Moses & Vardi 1995), while also throwing new light on old philosophical issues of verificationism (van Benthem 2004). As for a matching calculus of ‘hard information flow’, since these ideas are still less familiar than ‘hard core’ static epistemic logic, we state what the complete logic looks like:

Theorem *PAL* is axiomatized completely by the usual complete laws of epistemic logic plus the following *recursion axioms*:

$$\begin{aligned} [!P]q &\quad \Leftrightarrow \quad P \rightarrow q && \text{for atomic facts } q \\ [!P]\neg\phi &\quad \Leftrightarrow \quad P \rightarrow \neg[!P]\phi \end{aligned}$$

¹¹ Through the $!P$ and $[A]\phi$ steps, the two clauses of this definition involve a mutual recursion.

$$\begin{aligned} [!P]\phi \wedge \psi &\leftrightarrow [!P]\phi \wedge [!P]\psi \\ [!P]K_i\phi &\leftrightarrow P \rightarrow K_i(P \rightarrow [!P]\phi) \end{aligned}$$

while common knowledge reduces to ‘conditional common knowledge’:

$$[!P]C_G\phi \leftrightarrow (P \rightarrow C_G^P [!P]\phi) \quad ^{12}$$

These axioms are the ‘recursion equations’ of public information flow, performing step-by-step analysis of epistemic effects of incoming hard information. In particular, the final equivalence relates the knowledge that agents get after receiving new information to *conditional knowledge* they already had before.¹³ *PAL* is a simple system of what is arguably just the common sense view of semantic information. Even so, there is more to it than meets the eye, including a bisimulation-based model theory (van Benthem 2006B).¹⁴ Richer systems of *dynamic-epistemic logic* (*DEL*; cf. Baltag, Moss, Solecki 1998, van Benthem, van Eijck & Kooi 2006, van Ditmarsch, van der Hoek & Kooi 2007) deal with information flow in much more complex scenarios, such as *card games*, where not all players have equal observational access to events like drawing a card from the stack.

For the purposes of this paper, this further theory is beside the point. Instead, we now shift our focus to *consequence relations* naturally associated with *PAL*, to explore links with the pluralist program. Clearly, *PAL* has standard classical consequence, but it also has others. Thus, we can also look at it as generating new notions of consequence.

3.2 Structural rules for dynamic inference

Going back to the Restaurant example, here is a natural notion of consequence associated with what took place there. One first processes the information provided by the successive premises, and then checks the conclusion. To simplify, consider the propositional inference


‘from $A \vee B, \neg A$ to B ’,

¹² Van Benthem, van Eijck & Kooi 2006 define *P*-conditional common knowledge as common knowledge with accessibility restricted to finite paths consisting entirely of *P*-worlds, and provide the full-blown complete recursion axiom for it: $[!P]C_G^\phi \psi \leftrightarrow (P \rightarrow C_G^{P \wedge [!P]\phi} [!P]\psi)$.

¹³ Strictly speaking, this assumes *perfect memory* and other idealized epistemic features of agents.

¹⁴ Modern versions also include ‘*protocol information*’ about the total conversational or learning process that the individual updates are part of (van Benthem, Gerbrandy & Pacuit 2007).

starting from an initial situation where the agent has no knowledge about A and B . Here are the ‘updates’ for the two premises, ruling out 3 of the 4 options:



Next observe that updating with the conclusion B would not change the information state any more. Inspired by this rather natural scenario, one can define the following notion of abstract dynamic consequence (‘update-to-test’; Veltman 1997, van Benthem 1996):¹⁵

Definition A sequent $P_1, \dots, P_k \Rightarrow \phi$ is *dynamically valid* if, starting with any epistemic model (\mathbf{M}, s) whatsoever, successive announcements of the premises result in a model where announcement of ϕ effects no further change: i.e., in the model $(\dots(\mathbf{M}|P_1)\dots)|P_k, s)$ the formula ϕ was already true everywhere, even before it was announced.¹⁶ ♣

Modulo a few details, dynamic validity amounts to *PAL* validity of the following dynamic-epistemic formula, which says that the conclusion becomes *common knowledge*:

$$[!P_1] \dots [!P_k] C_G \phi \quad (\#) \quad ^{17}$$

In the case of a single S5-agent, which we will consider henceforth for convenience, we can replace the common knowledge modality $C_G \phi$ here by just $K\phi$. On the surface, this seems quite close to classical consequence. Indeed, the following is easy to see:

¹⁵ This notion is actually taught to students in Amsterdam, because it ‘feels right’ to them.

¹⁶ This notion is ‘partial’: it does not presuppose that all premises can be truthfully announced. Hans van Ditmarsch (p.c.) remarks that one might consider ‘local’ versions where we stop when the actual world satisfies the conclusion, and he has made some interesting observations about dynamic consequence involving this point – but my original examples call for the ‘global fixed-point’.

¹⁷ Here, validity refers to the *Supermodel* of all epistemic models related by arbitrary announcement steps. But when modeling realistic scenarios of conversation or enquiry, we can also relativize this to smaller *restricted families* \mathbf{M} of epistemic models, with *protocols* of admissible announcements.

Fact For purely factual (non-epistemic) formulas P_1, \dots, P_k, ϕ , dynamic consequence holds if and only if ϕ follows classically from P_1, \dots, P_k .

But the special reason why this holds is this: factual formulas do not change their truth values at worlds when passing from a model M to an updated model $M|P$. Things change when we admit announcements of epistemic formulas, because then, truth values can and will typically change, like when ignorance changes into knowledge.

Fact All classical structural rules fail for dynamic validity.


Proof We give a few cases, all using ‘Moore-type’ infelicities of the form $\neg Kp$ & p in making announcements, which lead to their own falsity. *Permutation* fails because $[!p][!<K>\neg p] K\perp$ is valid (the first announcement leaves only p -worlds; and so the second cannot be performed successfully), whereas $[!<K>\neg p] [!p] K\perp$ is not valid: the initial sequence of announcements is perfectly consistent, so \perp does not result necessarily. Likewise, *Contraction* fails since the update sequence $!<K>\neg p; !p$ is consistent, whereas the repeated $!<K>\neg p; !p; !<K>\neg p; !p$ is not. Finally, *Cut* fails as follows: we have $[!\neg p] K\neg p$ and $[!\neg K\neg p][!K\neg p] K\perp$, but we do not have $[!\neg K\neg p][!\neg p] K\perp$ valid. ♣

Now, precisely the same phenomenon emerges that we already know from general sub-structural logic. There are *modified structural rules* which do remain valid in this setting.

Fact Dynamic consequence satisfies the following structural rules:

- if $P \Rightarrow C$, then $A, P \Rightarrow C$ *Left-Monotonicity*
- if $P \Rightarrow A$ and $P, A, Q \Rightarrow C$, then $P, Q \Rightarrow C$ *Left-Cut*
- if $P \Rightarrow A$ and $P, Q \Rightarrow C$, then $P, A, Q \Rightarrow C$ *Cautious Monotonicity*

Now we are at the abstraction level of structural rules, and indeed, these rules are valid in a much more general setting. We can view propositions A dynamically as *partial functions* T_A taking input states meeting the preconditions of update with A to output states:



Definition Abstract transition models $\mathbf{M} = (S, \{T_A\}_{A \in Prop})$ consist of states S with a family of transition relations T_A for each abstract proposition A . Here, a sequence of propositions $\mathbf{P} =$

P_1, \dots, P_k dynamically implies conclusion C in \mathbf{M} , if any sequence of premise updates starting anywhere in \mathbf{M} ends in a fixed point for the conclusion: if $s_1 T_{p_1} s_2 \dots T_{p_k} s_{k+1}$, then $s_{k+1} C s_{k+1}$. We say sequent $P_1, \dots, P_k \Rightarrow C$ is true in the model: $\mathbf{M} \models P_1, \dots, P_k \Rightarrow C$. ♣

It is easy to check that the above three structural rules hold even for this abstract setting. Moreover, van Benthem 1996, Chapter 7, proves the following representation result:

Theorem A sequent σ is derivable from a set of sequents X by these three rules iff σ is true in all models where all sequents in X are true.

The argument is a neat syntactic construction whose details we forego here.¹⁸ This abstract analysis of a natural ‘dynamic’ sub-structural consequence relation seems to extract the ‘gist’ of inference in dynamic-epistemic logic. But to show that it really does, we need to tighten up the connection. Here is a sketch of how this can be done (van Benthem 2003B has the details). First we need to introduce the following generalized notion:

Definition A meta-sequent $\Sigma \Rightarrow \sigma$ from a set of sequents Σ to a sequent σ is *update-valid* if all its substitution instances with epistemic formulas, reading sequents like before as type (#) dynamic-epistemic formulas, gives a valid implication between *PAL*-formulas.¹⁹ ♣

Theorem The update-valid structural inferences $\Sigma \Rightarrow \sigma$ are precisely those whose conclusions σ are derivable from their premise sets Σ by the rules of Left-Monotonicity, Left-Cut, and Cautious Monotonicity.

Proof Soundness is by immediate inspection. Completeness uses two representation steps. One finds a counterexample on an abstract transition model as above, and then transforms this into a concrete family of epistemic models for the states,²⁰ and concrete announcement actions for the labeled transitions (the construction is in van Benthem 2003B). ♣

¹⁸ These representations involve a sort of hunt for ‘poor man’s completeness theorems’.

¹⁹ For the special ‘universal Horn’ formulas obtained in this way, validity in the above Supermodel, or in arbitrary more constrained ‘protocol models’ \mathbf{M} as above, makes no difference.

²⁰ The transformation works up to modal bisimulation, to be a bit more precise.

3.4 Natural levels of abstraction: modal logic of dynamic consequence

But actually, the leap from *PAL* to the preceding austere sequent-style analysis is rather drastic. There are other natural abstraction levels behind our *PAL*-style dynamic-epistemic logic of agency. In particular, there is no reason whatsoever why an insightful analysis of a notion of consequence could not have some well-chosen operators in its language!

Indeed, the above transition models are really just models for a standard *poly-modal logic*. The above notion of dynamic validity needs *two* basic kinds of modality, viz. (a) universal modal boxes for the premise transitions, and (b) a ‘loop modality’ for the fixed-points:

$$\mathbf{M}, s \models (a)\phi \text{ iff } R_a ss \ \& \ \mathbf{M}, s \models \phi^{21}$$

The modal loop language is decidable, and its complete axiomatization has key axioms

$$(a)\phi \leftrightarrow (a)T \ \& \ \phi, \quad (a)T \rightarrow ([a]\phi \leftrightarrow \phi).$$

Reading dynamic sequents $P_1, \dots, P_k \Rightarrow C$ as modal formulas $[P_1] \dots [P_k](C)T$,²² all earlier structural rules become very simply derivable in this language. But the modal language can also express complex existential properties of consequence beyond mere structural rules. Thus, poly-modal logic seems a natural stage for a richer abstract theory of dynamic inference. Moreover, it still stays close to the original setting of dynamic-epistemic logic, as may be seen by extending our earlier definitions for the case of mere sequents:

Theorem The *update-valid modal formulas* are axiomatized precisely by the general minimal modal logic of $[a]$ and (a) for partial functions a .

The proof in van Benthem 2003B uses a representation of arbitrary finite modal tree models into *PAL*-universes with announcements running between models.

There are still further natural abstraction levels than the two considered in this section. For instance, going just one step beyond *PAL*, we can ask for the *schematic validities* of the


²¹ Added to *PAL*, such fixed-point operators add expressive power: cf. Baltag & Smets 2007.

²² Note the difference with our earlier treatment of conclusions in terms of common knowledge.

system which remain valid whatever we substitute for their proposition letters.²³ It is not known if this is decidable or even axiomatizable (van Benthem 2006B).

3.5 Conclusion

We have shown that the relation between abstract analysis of consequence relations and dynamic logics of agency can be interesting and fun, giving rise to non-trivial questions. One can look at this connection in two directions, as in the following schema:



‘From dynamics to consequence’, the issue is to find good *abstraction levels* capturing significant properties of consequence relations generated by the concrete activity modeled in the dynamic logic. ‘From consequence to dynamics’, one reconstructs (or just brings out of the closet) the dynamic practice generating the given consequence relation, and this is what representation theorems do. The two directions obviously live in harmony, and we can perform a Gestalt Switch one way or the other. The two directions might also be used to describe historical periods. The avant garde tendency in applied areas in the 1980s was toward abstraction, and maybe that of the current decade more towards concretization.

4 Non-monotonic reasoning and dynamic logic of belief change

With this case study of knowledge and cumulative information update in place, let us now move to the next level of agency in our earlier ladder, which involves ‘jumps’ in the form of actions of belief revision, self-correction, and learning. This makes another major branch of substructural logic a test case, viz. *non-monotonic logics* generated by default reasoning, circumscription, and the like. This time, we reverse the perspective chosen in Section 3. We start with abstract formats for non-monotonic reasoning, and their dynamification in dynamic-epistemic style comes only afterwards. Finally, we draw comparisons again.

4.1 Minimization over preference orders

Classical logical consequence from premises P to conclusion C says all models of P are models for C . The famous insight in McCarthy 1980 was that human and machine problem

²³ Note that the stated reduction axiom for atomic propositions is not valid in this sense.

solving and planning go beyond this, getting more out of premises by zooming in on the most 'congenial' models. A *circumscriptive* consequence from P to C says that

C is true in all the *minimal* models for P

Here, minimality is taken with respect to some relevant comparison order \leq for models: inclusion of object domains, inclusion of denotations for specified predicates, and so on. The general idea is minimization over any reflexive transitive order of 'relative plausibility' (Shoham 1988), much as in the Lewis semantics for *conditional logic* since around 1970 – an analogy often noted (cf. Gaerdenfors & Rott 1995). One can study these consequence relations in terms of structural rules, as has been done by many authors following Gabbay (cf. Gabbay 1996).²⁴ But conditional logic *itself* is also an interesting candidate for a natural abstraction level, adding Boolean connectives, and it may be compared to the above modal logic in describing consequences. We merely cite one perhaps less-known extremely simple representation result to show the spirit of working at this level (van Benthem 1989).

Update by minimal elements Consider any set E with a binary order \leq . First the map f which takes a set X to the subset $\min(X)$ of all \leq -minimal elements in X : $\{x \in X \mid \forall y \in X: x \leq y\}$ satisfies the following three set-theoretic conditions:

- $C1$ $f(X) \subseteq X$
- $C3$ $f(X) \cap Y \subseteq f(X \cap Y)$
- $C4$ $\bigcap_{i \in I} f(X_i) \subseteq f(\bigcup_{i \in I} X_i)$

Indeed we have an equivalence here:

Fact The following conditions are equivalent:

- (a) f satisfies $C1$, $C3$, $C4$
- (b) there is an order \leq on E s.t. $f(X) = \min(X)$ for all X

Proof From (a) to (b). Define a binary order as follows: $x \leq y$ iff $x \in f(\{x, y\})$.

Claim $f(X) = \min(X)$

²⁴ Cf. various chapters in the *Handbook of Logic in Artificial Intelligence and Logic Programming*.

Proof The inclusion from left to right. Let $x \in f(X)$: then $x \in X$ by *C1*. Let $y \in X$. Now in condition *C2*, take $Y = \{x, y\}$. Then we have $x \in f(X) \cap Y \subseteq f(X \cap Y) = f(Y)$, and hence that $x \leq y$. Next, the inclusion from right to left \supseteq . Let $x \in \min(X)$: i.e., for all $y \in X$, $x \in f(\{x, y\})$. Then x is in the intersection of all these sets $f(\{x, y\})$. By *C3*, then, $x \in f(\bigcup_{y \in X} \{x, y\})$ ♣

Likewise, belief revision theory started at this austere level in terms of the *AGM*-postulates (Gaerdenfors 1987), which merely constrain concrete rules for changing one's mind. The resulting theory of all this is well-known, and it seems to support a consequence-based perspective, since belief revision theory is often equated with non-monotonic logic.

Let us now move upward toward our Logical Dynamics program, and shake the tree a bit. I would like to suggest that a shift in perspective may be helpful – from a steaming jungle of non-classical 'consequence relations' to the current world of modal logics for belief update, belief revision, and other *informational attitudes* and *informational processes*.

4.2 Dynamification: non-monotonic logic as monotonic logic of belief revision

Let us return to the puzzles that motivated non-monotonic logic in the first place. We are given some initial information, and need to find out the true situation. Extra information may come on the way. I submit that the most striking phenomenon in such scenarios is *not inference at all*, but rather our receiving that information, and our subsequent responses:

We are playing the board game "Kings and Cardinals" (the board is an object of public observation) having 'monasteries' and 'advisors' placed here and there. I look at the cards in my hand (a private observation), and also at the map of medieval Europe on the board. Right now, I know certain things about the outcome of the game, and I believe more than what I strictly know, based on my expectations about cards that the other players hold, or their temperaments: timid, bluffing,... Now, new information comes in: you select a new country on the map and place some counters there. This observation changes my current information state. I know more now, and the observation may even speed along further beliefs of mine: you are trying to build a trade route from Burgundy to Bohemia. Of course, these current beliefs may be refuted by further moves of yours, unlike the hard infeasible knowledge which I have obtained about what's on the board.

Solving puzzles and playing games is all about such processes. But this is precisely the arena of our *dynamic logics of information update*, provided we can also make them deal

with belief revision. The very motivation for non-monotonic reasoning seems epistemic or doxastic, having to do with managing knowledge and beliefs – but this key feature *is left implicit*. We have already seen that classical consequence is about the *knowledge update* that takes place when new information comes in. And in tandem with this, I would say that

Circumscriptive inference is about belief formation

which takes place on the basis of incoming new information. Clearly, knowledge update and belief revision are intertwined, and they provide mutual support. I think it is this diversity of responses to information which truly explains the modern galaxies of 'notions of consequence', where different styles live together. So, let's look at some modal logics underpinning these phenomena when we shift the focus to processing information.

4.3 Dynamic logic of belief change under hard information

I start with the logic which is closest to minimizing non-monotonic logic, though it may not be the most interesting one in the end. But we first need to get some statics in place.

Belief and plausibility order Agents have other attitudes to propositions than knowledge, in particular, *beliefs* that may turn out incorrect. Logics of belief analyze assertions

$B_i\phi$ for 'agent i believes that ϕ '.

Their semantics adds further gradations to the information ranges in epistemic modeling, in the form of a *plausibility ordering* of worlds x, y as seen from some vantage point s :

$\leq_{i,s} xy$ in world s , agent i considers y at least as plausible as x .

In particular, we now define belief semantically as '*truth in the most plausible options*':

$\mathbf{M}, s \models B_i\phi$ iff $\mathbf{M}, t \models \phi$ for all t which are maximal in the ordering $\lambda xy. \leq_{i,s} xy$.²⁵

There are some complications with making this stipulation in infinite models, but what we have here is the main idea. Incidentally, what we really have in mind is that we look at the most plausible worlds among those that are *epistemically accessible* to the current world. We suppress this fact notationally, but it may be useful to keep it in mind when reading on.

²⁵ Of course, working with 'maximality' is just as good as our earlier use of minimality.

Example Consider a model with two possible worlds that are mutually epistemically accessible, but the one with $\neg P$ is considered more plausible than the other:

$$\boxed{P \text{ --- } \leq \text{ --- } \neg P}$$

At the actual world with P , the agent does not know whether P , but she does (mistakenly!) believe that $\neg P$. It is crucial that our beliefs can be false.

For complete doxastic logics and more theory around them, cf. Fagin et al. 1995.²⁶

Next, in doxastic logic, one soon finds that absolute beliefs are not sufficient for explaining agents' behaviour. We want to know what they would believe were they to receive new information. This *pre-encoding*, in our earlier sense, requires *conditional belief*:

$$\begin{aligned} \mathbf{M}, s \models B_i^\psi \phi \quad \text{iff} \quad \mathbf{M}, t \models \phi \text{ for all worlds } t \text{ which} \\ \text{are maximal for } \lambda xy. \leq_{i,s} xy \text{ in the set } \{u \mid \mathbf{M}, u \models \psi\}. \end{aligned}$$

Conditional beliefs $B_i^\psi \phi$ are again like general conditionals (cf. Lewis 1973), in that they express what might happen under different circumstances from where we are now.²⁷

Dynamic-doxastic logic of hard information Now, consider the effect of the earlier events of receiving public hard information on agents' *beliefs*. Given the availability of conditional beliefs, we get this result from van Benthem 2007A, comparable to our earlier *PAL*:

Theorem The complete logic of conditional belief under public announcements is axiomatized by (a) any complete static logic for knowledge and belief, (b) the *PAL* reduction axioms for atomic facts and Boolean operations, (c) the following new recursion axiom for conditional beliefs:

$$[!P] B_i^\psi \phi \leftrightarrow (P \rightarrow B_i^{P \wedge [!P]\psi} [!P]\phi) \quad ^{28}$$

²⁶ Most logics also analyze the *interplay* between knowledge and belief in information models with two relations \sim_i, \leq_j entangled in various ways, reflecting a stand on whether knowledge implies belief, or whether one knows one's beliefs. While relations between attitudes toward information are an important topic, we focus on belief in what follows.

²⁷ The analogy is so close that conditional belief on reflexive transitive plausibility models satisfies exactly the laws of the minimal conditional logic (cf. Veltman 1985).

There is more to this system than meets the eye. For instance, the stated reduction axiom does not straightforwardly reduce beliefs after update to conditional beliefs one had before (a popular view): this will only be the case when all propositions involved are factual. Moreover, scenarios of belief change under update can be tricky. As has been observed in both computer science and philosophy (Shoham and Leyton-Brown 2008, Stalnaker 1996), true information can trick agents from true into false beliefs, by ruling out most plausible but non-actual worlds where some true proposition holds. Thus, there is room for new notions in between knowledge and belief, such as ‘safe belief’ (truth in all worlds at least as plausible as the current one; Baltag & Smets 2007), which has its own recursion laws in the above style. The present formalism allows us to experiment with many doxastic notions.

4.4 Dynamic consequence and circumscriptive inference

The epistemic-doxastic setting suggests two notions of dynamic consequence – in the same style as in the preceding section. Both state what happens once the premises are processed: either knowledge results just as before²⁹ – or we go to belief:

$$[!P_1] \dots [!P_k] B\varphi$$

Given our semantics of belief, and working with factual propositions, the latter is precisely the dynamic counterpart to minimizing consequence relations like circumscription:

$$P_1, \dots, P_k \Rightarrow \varphi$$

Indeed, our claim is simply this. Circumscription leads to *beliefs* rather than knowledge, since its conclusions may be retracted on the basis of further evidence. But then, what has traditionally been cast as a new ‘non-standard’ consequence relation may also be seen differently through ‘dynamification’. Making the dynamic setting more explicit, we have a *dynamic logic of belief formation under incoming factual propositions*. Technically, the hall-mark failure of monotonicity then occurs because of the minimization in the definition of belief – not because of some special feature of the notion of consequence as such.

²⁸ As a special case, this describes formation of absolute beliefs: $[!P] B_i\phi \leftrightarrow (P \rightarrow B_i^P [!P]\phi)$.

²⁹ This should be *common knowledge* in the multi-agent case, and likewise later *common belief*.

But this setting is also richer than the consequence view. (a) The dynamic epistemic and doxastic language allows complex *non-factual* propositions for premises and conclusions. The usual accounts of structural rules and the ‘intuitions’ associated with them do not seem to take this option of more sophisticated information into account. But the above calculus will keep all this absolutely straight: we do have the complete logic for the total language! (b) Also, as we just observed, the semantic framework allows for new operators beyond knowledge and belief, like ‘safe belief’. And so it raises new questions. What would be the *new consequence relation* associated with using safe belief rather than plain belief? Finally, (c), in our next section, we discuss another dynamic degree of freedom in setting up consequence relations: *the way in which we add* the information from the premises.

4.5 Dynamic logic of belief change under soft information

Soft information and plausibility change In applications of circumscription, one fixes a comparison relation between models, which does not change in the process of inference. In abstract non-monotonic logics, this choice is even left implicit in context, without a trace in the formal language. But since comparing worlds by relative plausibility determines agents' beliefs, and hence their 'conclusions', it seems important to have explicit control over how we choose, and *change*, that ordering. Indeed, triggers for changing beliefs need not be 'hard information' of the public announcement type, ruling out certain worlds for good. They can rather be 'soft information' affecting just our plausibility ordering of the worlds!

A triggering event which makes us believe that P need only *rearrange worlds* making the most plausible ones P : it works by 'promotion' rather than elimination of worlds. Thus, on the earlier models $\mathbf{M} = (W, \sim_i, \leq_i, V)$, we change the relations \leq_i , rather than the domain of worlds W or the epistemic accessibilities \sim_i . Here is a well-known soft trigger from the area of belief revision, sometimes called 'radical revision'. A *lexicographic upgrade* $\uparrow P$ is an instruction for changing the current ordering relation \leq between worlds as follows:

all P -worlds in the current model become better than all $\neg P$ -worlds,
while, within those two zones, the old plausibility ordering remains.

We have the following corresponding dynamic modality

$$\mathbf{M}, s \models [\uparrow P]\phi \quad \text{iff} \quad \mathbf{M} \uparrow P, s \models \phi$$

with $\mathbf{M} \uparrow P$ the model \mathbf{M} with its order \leq changed as stated above. This dynamic doxastic language describes how beliefs change under soft information (van Benthem 2007A):

Theorem The dynamic logic of lexicographic upgrade is axiomatized completely by

- (a) any complete axiom system for conditional belief on the static models, plus
- (b) the following recursion axioms:

$$\begin{aligned} [\uparrow P]q &\leftrightarrow q, && \text{for all atomic proposition letters } q \\ [\uparrow P]\neg\phi &\leftrightarrow \neg[\uparrow P]\phi \\ [\uparrow P](\phi \wedge \psi) &\leftrightarrow [\uparrow P]\phi \wedge [\uparrow P]\psi \\ [\uparrow P]B^\psi\phi &\leftrightarrow (E(P \wedge [\uparrow P]\psi) \wedge B^{P \wedge [\uparrow P]\psi} [\uparrow P]\phi) \\ &\quad \vee (\neg E(P \wedge [\uparrow P]\psi) \wedge B^{[\uparrow P]\psi} [\uparrow P]\phi) \end{aligned} \quad ^{30}$$

The final equivalence describes which conditional beliefs agents form after soft upgrade. This may look daunting, but try to read the principles of some default logics existing today! And there is a reward. We now see explicitly how new triggers affect the plausibility order \leq among worlds, and hence our beliefs at any given stage, and thus, the 'nonmonotonic inferences' available to us on the basis of the ambient order \leq . Moreover, the formulation in terms of conditional belief at once solves the 'iteration problem' which has plagued belief revision: we do not just know the new beliefs, but the above axioms also completely describe the new tendencies that agents have toward further belief revision. ³¹

Other options In addition, there are many further possible ways of taking soft information. For instance, a more conservative form of belief revision puts not *all* P -worlds on top qua plausibility, but just *the most plausible* P -worlds. 'After the revolution', this policy co-opts

³⁰ Here, ' E ' is either a global existential modality, or the epistemic existential modality $\langle K \rangle$.

³¹ I am not clear on one thing here: the above consequence relations do not use conditional beliefs in their conclusions. But this may just reflect expressive poverty of the format. If you want rules of *conditionalization* appropriate to your setting, then conditional operators have to come in after all.

just the leaders of the underclass, not all of them. Complete dynamic logics for these and other policies exist, too (van Benthem 2007A; Baltag & Smets 2007).

4.6 Variations on circumscriptive consequence relations

Given these options in belief revision and plausibility change, which is the true analogue of a circumscriptive or general ‘minimizing’ inference $P \Rightarrow \phi$ in this dynamic setting? In particular, intuitively, are the premises instances of hard information or soft information? We cannot tell, because we now live in a richer universe of informational events that may determine how we solve our problem, make our plan, or play our game.

The two options will not be the same, of course, and their different behaviour is described by our complete dynamic doxastic logics. Even so, there is an interesting type of question, similar to the one raised in our section about information update and knowledge change:

What are complete sets of structural rules for the consequence relations:

$$P_1, \dots, P_k \Rightarrow_{\text{circ-hard}} \varphi \quad \text{iff} \quad [!P_1] \dots [!P_k] B\varphi$$

$$P_1, \dots, P_k \Rightarrow_{\text{circ-soft}} \varphi \quad \text{iff} \quad [\uparrow P_1] \dots [\uparrow P_k] B\varphi$$

I have no answer, though I expect the theory of the first is a mix of the usual sub-structural rules for non-monotonic logic plus the additional dynamic phenomena in Section 3. But here is at least a structural difference between the two notions, even for factual assertions:

Fact For factual assertions P, Q , (i) $P, Q \Rightarrow_{\text{circ-hard}} P$, (ii) $\text{not } P, Q \Rightarrow_{\text{circ-soft}} P$.

Proof (i) Successive hard updates yield subsets of the P -worlds. (ii) The last upgrade with Q may have demoted all P -worlds from their former top positions. ♣

How to choose between such alternative notions? It all depends on the scenario of problem solving or game playing that we are engaged in. Indeed, our logic provides many more, once we look at other plausibility-changing events. And the shifts in that plausibility order are really the primary issue in understanding how we navigate through the task at hand.

4.7 Summary

The intuitions behind circumscriptive inference styles involve knowledge and belief. They are also dynamic, involving agents’ responses to incoming information. Thus, in a dynamic

epistemic perspective, circumscription and other styles of non-monotonic reasoning are at heart about cognitive attitudes and responses to information. Moreover, these responses can be quite different, from hard information update to soft plausibility change. Merging things in this way fits with the general conception of agency stated in Section 2: processes of *inference* and *self-correction* go hand in hand! This perspective also generates new consequence relations, and new employment for sub-structural analysis. We see this as validation of the perspective advocated in Section 3: what we saw with update and classical consequence becomes even more interesting with revision and non-monotonicity.³²

5 Conclusion and outlook

Pluralism meets Dynamics This paper has contrasted two programs for legitimizing the diversity of modern logic. ‘Logical Pluralism’ emphasizes consequence as the locus of research, and finding natural ways to parametrize it. Its ‘strong arm’ is mathematical analysis of possible consequence relations. The other program is ‘Logical Dynamics’, emphasizing events of information flow, from inference to observation, and the various processes by which rational agents harness this to act and interact. Its formal paradigm is dynamic logic in some suitably broad sense. We have shown how one can move back and forth between the two perspectives by processes of abstraction and ‘dynamification’, and sometimes even get very precise connecting results (after all, we are in the same field). Much more can happen here, once we put the research agendas in the two programs side by side.³³ But, I am not completely neutral in my evaluation. Even though I started out in the 1980s on the sub-structural consequence side, I would now prefer the dynamic perspective – partly because it is more ambitious, and gives logicians many more things to do!

The case of circumscription More specifically, we have shown in a specific case study of circumscriptive-style minimizing inferences, how it might be legitimate to question the

³² Of course, it would add support to my general position in this paper if we could do such a dynamified reconstruction of other non-monotonic reasoning styles, too: say, for *abduction*.

³³ This also raises issues. What are natural levels of abstraction for a given information-handling activity? And, is not dynamification, which makes ever more things explicit in the logical language well on its way toward *trivialization*, viz. translation into some suitably rich meta-language?

original immediate impulse of McCarthy and his many followers to cast things as a notion of consequence, rather than a process of belief revision. What we saw in Section 4 is that, if we do, a non-monotonic consequence relation may also be seen as a classical dynamic logic of the process which *causes* the non-monotonicity. In a slogan, *monotonic dynamic logic can model non-monotonic consequence!* The general scope of this slogan remains to be understood, and it depends of course on successful further cases of dynamification.

Challenges to dynamics Actual research in the Logical Dynamics tradition also poses a number of challenges to the dynamic logic approach. Perhaps the most obvious one is the work on relevant logic, resource-sensitive linear logics, situation-theoretic logics, and other systems that deal with both information structure and information flow. These all involve a process which we have ignored here, viz. *inferential dynamics*. And thus, they raise a largely unresolved issue of reconciling the different notions of information that play in modern logic (van Benthem & Martinez 2007). This is all the more pressing since my own earlier work on categorical grammars and information structure was squarely in that tradition (van Benthem 2008C is a first attempt at putting together the pieces of my life). I have nothing of substance to say on this here, but it is a major issue to be resolved.³⁴

Further challenges to the dynamics program include *para-consistent logic* and making its underlying processes of inconsistency handling explicit, as well as more ambitious issues of *language change* (reflecting conceptual changes) in response to incoming events. The latter interest would go back to Bolzano after all, who did include the choice of language as an explicit and crucial parameter in his account of logical consequence.

³⁴ For instance, the dynamic consequence behind the categorical Lambek Calculus (van Benthem 1991) is this. Let propositions A be any *transition relations* R_A between abstract states, and say that P_1, \dots, P_k *dynamically implies* conclusion C in a model, if any sequence of premise updates starting anywhere in effects a total transition for the conclusion: *if* $s_1 R_{p_1} s_2 \dots R_{p_k} s_{k+1}$, *then* $s_1 R_C s_{k+1}$. (contrast this with the dynamic consequence of Section 3). Interpretations of these transitions range from syntactic concatenation to abstract information merge. Comparing this with the dynamic-epistemic view of events that trigger information flow is feasible (I have some first technical results), but it quickly leads to divergences from the bisimulation-based modal framework.

Once again: what is logic? The background to this paper is the question what logic really is. Some seek this in mathematical notions of semantic invariance (Bonney 2006), others in some proof-theoretic minimum of structural rules that should never be violated (cf. Martin–Löf 1996 for a sophisticated view of proof-theoretic foundational perspectives). To me, even Logical Pluralism stays close to the traditional foundations of mathematics, since it sees logic as being about consequence relations and formal systems.³⁵ But I myself see much of our current discussion as trying to break away from the magnetic spell of those mind grooves formed in the grand foundational period of the 1930s. My views on this were stated in Section 2, including a turn from proof and computation as paradigms toward rational agency and intelligent interaction.³⁶ I see the nature of logic in its features as a dynamic activity, not as any static product of that activity: proofs, formal systems, or languages. And that is Logical Dynamics again, now applied to the whole field.

References

- 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.
- A. Baltag, L. Moss & S. Solecki, 1998, ‘The Logic of Public Announcements, Common Knowledge and Private Suspicions’, *Proceedings TARK 1998*, 43 – 56, Morgan Kaufmann Publishers, Los Altos.
- A. Baltag & S. Smets, 2007, ‘Dynamic Belief Revision over Multi-Agent Plausibility Models’, Computing Lab, Oxford University. An earlier version appeared in *Proceedings LOFT 2006*, Department of Computing, University of Liverpool.
- J. Beall & G. Restall, 2006, *Logical Pluralism*, Oxford University Press, Oxford.

³⁵ Even the wonderful modernist volume Gabbay 1993 channeled all modern discussion of this sort into the mathematical issue of ‘What is a Logical System’?

³⁶ I believe this move toward agency is also the current view of Dov Gabbay, though of course, ‘anything I say in this paper is my sole responsibility’, and I am happy to fall alone...

- J. van Benthem, 1985, 'The Variety of Consequence, According to Bolzano',
Studia Logica 44:4, 389 – 403.
- 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, 1996A, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- J. van Benthem, 2003A, 'Is there still Logic in Bolzano's Key?', in E. Morscher, ed.,
Bernard Bolzanos Leistungen in Logik, Mathematik und Physik, Bd.16,
Academia Verlag, Sankt Augustin 2003, 11 – 34.
- J. van Benthem, 2003B, 'Structural Properties of Dynamic Reasoning', in J. Peregrin,
ed., *Meaning: the Dynamic Turn*, Elsevier, Amsterdam, 15 – 31.
- J. van Benthem, 2004, 'What One May Come to Know', *Analysis* 64 (282), 95–105.
- 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, 2007A, 'Dynamic Logic of Belief Revision', *Journal of Applied
Non-Classical Logics* 17, 129 – 155.
- J. van Benthem, 2007B, 'Logic, Rational Agency, and Intelligent Interaction', invited
lecture at the 13th International Congress of Logic, Methodology, and
Philosophy of Science, Beijing 2007, to appear in D. Westerståhl et al.,
eds., *Proceedings DLMPS XIII*, College Publications, London.
- J. van Benthem, 2008A, 'Merging Observation and Access in Dynamic Logic',
Studies in Logic, 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, 2008C, ‘The Information in Intuitionistic Logic’, Department of Philosophy, Stanford University. To appear in *Synthese* issue on ‘Philosophy of Information’ edited by Luciano Floridi & Sebastian Sequoiah-Grayson (Oxford).
- J. van Benthem, J. van Eijck & B. Kooi, 2006, ‘Logics of Communication and Change’, *Information and Computation* 204(11), 1620 – 1662.
- J. van Benthem, J. Gerbrandy & E. Pacuit, 2007, ‘Merging Frameworks for Interaction: *DEL* and *ETL*’, ILLC Amsterdam & Informatics Torino. *Proceedings TARK 2007*, University of Namur. Extended version 2008, with Tomohiro Hoshi, 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, R. Muskens & A. Visser, 1997, ‘Dynamics’, in J. van Benthem & A. ter Meulen, eds., *Handbook of Logic and Language*, Elsevier, Amsterdam, 589 – 647.
- B. Bolzano, 1837, *Wissenschaftslehre*, Seidelsche Buchhandlung, Sulzbach. Translated as *Theory of Science* by R. George, University of California Press, Berkeley & Los Angeles, 1972.
- D. Bonnay, 2006, *What is a Logical Constant?*, Ph.D. thesis, École Normale Supérieure, Paris.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, 2007, *Dynamic Epistemic Logic*, Springer, Dordrecht.
- M. Dunn, 1991, ‘Gaggle Theory: An abstraction of Galois connections and Residuation, with applications to negation, implication, and various logical operators’, in J. van Eijck, ed., *Logics in AI (Amsterdam, 1990)*, Springer, Berlin, 31 – 51.
- K. Dosen & P. Schroeder-Heister, eds., 1993, *Substructural Logics*, Oxford University Press, Oxford.
- R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, The MIT Press, Cambridge (Mass.).
- D. Gabbay, ed., 1993, *What is a Logical System?*, Oxford University Press, Oxford.

- D. Gabbay, 1996, *Labeled Deductive Systems*, Oxford University Press, Oxford.
- P. Gärdenfors, 1987, *Knowledge in Flux, on the Dynamics of Epistemic States*, The MIT Press, Cambridge (Mass.).
- P. Gärdenfors & H. Rott, 1995, 'Belief Revision', in D. M. Gabbay, C. J. Hogger & J. A. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming 4*, Oxford University Press, Oxford 1995.
- D. Lewis, 1973, *Counterfactuals*, Blackwell, Oxford.
- P. Lorenzen, 1955, *Einführung in die Operative Logik und Mathematik*, Springer, Berlin.
- P. Martin-Löf, 1996, 'On the Meanings of the Logical Constants and the Justifications of the Logical Laws', *Nordic Journal of Philosophical Logic*, 1, 11 – 60.
- J. McCarthy, 1980, 'Circumscription – A Form of Non-Monotonic Reasoning', *Artificial Intelligence* 13, 27 – 39.
- G. Restall, 2000, *An Introduction to Substructural Logics*, Routledge, London.
- D. Scott, 1971, 'On Engendering an Illusion of Understanding', *The Journal of Philosophy*, Vol. 68, No. 21, 787 – 807.
- Y. Shoham, 1988, *Reasoning about Change*, The MIT Press, Cambridge (Mass.).
- Y. Shoham and K. Leyton-Brown, 2008, *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*, Cambridge University Press, Cambridge.
- R. Stalnaker, 1996, 'Knowledge, Belief and Counterfactual Reasoning in Games', *Economics and Philosophy*, 12, 133 – 163.
- F. Veltman, 1985, *Logics for Conditionals*, Dissertation, Philosophical Institute, University of Amsterdam.
- F. Veltman, 1997, 'Defaults in Update Semantics', *Journal of Philosophical Logic*, 25, 221 – 261.