Journal of Logic and Computation Advance Access published July 7, 2014

Dynamic consequence for soft information

OLIVIER ROY, *Department of Philosophy, Universität Bayreuth, Germany*
E-mail: Olivier.Roy@uni-bayreuth.de

OLE THOMASSEN HJORTLAND, *Chair of Logic and Philosophy of Language, LMU Munich, Germany*

Abstract

This article looks at so-called dynamic consequence relations for models of soft information change. We provide a sound, complete calculus for one-step soft dynamic consequence relations. We then study a generalization to sequences of updates, for which we show a number of valid and invalid structural rules.

Keywords: Dynamic epistemic logic, substructural logic, soft information.

Dynamic consequence relations are consequence relations generated by concrete information update procedures. Here we study soft information updates, which are updates that are reversible and not necessarily truthful. We provide a sound and complete axiomatization for dynamic consequence relations generated by a large class of such updates. The calculus covers one-shot updates. It describes the effect of a ψ-type of update in contexts where ϕ holds. Towards the end of the article we generalize this calculus to cover sequences of information updates, also known as iterated revisions. We show a number of valid inferences in that calculus, and use them to compare dynamic and classical consequence relations.

1 Why dynamic consequence relations

Rational belief change can be seen as licensing specific kinds of inferences, that is, inducing a specific kind of consequence relation. In the words of [12], belief change and non-classical inferences are 'two sides of the same coin'. This idea, by now widely accepted in default logic, also underlies the notion of *dynamic consequence*.

This article focuses on one specific framework for the study of information change in multi-agent settings: *dynamic epistemic logic* (DEL). See [26] and [22] for an in-depth presentation. The gist of DEL is to enrich propositional modal languages with operators that describe the effects or consequences of certain 'epistemic actions' in a given situation. Epistemic actions are events that only affect the information available in a given situation, such as observing certain states of affairs or learning about them from a truthful and trusted source.[1]

Dynamic consequences for DEL were first proposed by [21] in the context of *public announcement logic*. Public announcement logic contains formula $[!\phi]\psi$, to be read 'ψ holds after the announcement of ϕ'. One can use these formulas to define a consequence relation as follows:

$$\phi_1,...,\phi_n \models_{dyn} \psi \text{ iff } \mathcal{M},w \models [!\phi_1]...[!\phi_n]\psi \text{ for all pointed models } \mathcal{M},w.$$

[1] Epistemic actions are usually distinguished from 'ontic events', i.e. events, like turning on the lights, which change some non-informational facts. Ontic events have also been studied in DEL, cf. [23]. We do not study them here.

doi:10.1093/logcom/exu040

Dynamic consequence relations are interesting for a number of reasons. First, they are sub-structural. van Benthem [21] shows for instance that structural rules like cut, contraction and weakening fail for the relation just defined. We see in Section 5 that this also holds true for dynamic consequence relations for soft information changes. Such consequence relations therefore fall within the broad family of sub-structural logic [19]. It is well known that these logics can be given an abstract 'informational interpretation', and that this interpretation can be used to argue for or against certain structural rules—see e.g. [16]. But the strength of such arguments depends on the plausibility of the notion of information and information flow that are used in interpreting substructural systems. This can be hard to assess in the abstract. Dynamic consequence circumvents this difficulty by arguing the other way around. It starts with a concrete, well-understood belief update procedure, and uses this to construct a consequence relation that turns out to be sub-structural. For dynamic consequence relations the motivation to 'go sub-structural' comes from the plausibility of the underlying belief update mechanism.

Dynamic consequence relations are also contributing to our understanding of information change in social interaction. The standard technique to axiomatize validities in DEL is to use so-called *reduction axioms*, which show us how to analyse, compositionally, the effects of epistemic action in terms of what holds in the initial situation.[2] Recently, cut-free, complete Gentzen-style systems for DEL have also been developed, using labelled sequents or display logic [2, 13, 14]. These approaches say little about the valid logical operations on the epistemic actions themselves. These operations are conceptually important. Different epistemic actions are different types of learning events,[3] for example a public announcement that ϕ to a group *vs* a private announcement of ϕ to all members of that group. Valid and invalid inferences about epistemic actions show how these different learning events relate to one another. Structural rules in dynamic consequences make these inferences explicit, and thus unveil logical principles governing the dynamic of information that were hitherto implicit or not well studied.

Finally, echoing Cordón-Franco *et al*. [8], it should be noted that dynamic consequence relations can be used for reasoning about protocols in multi-agent systems and that they put DEL in a historical perspective. They indeed indicate that dynamic consequence relations constitute a return to the original motivation behind *dynamic semantics* [20, 27], which arguably lies at the source of DEL.

2 State of the art and present contribution

The study of dynamic consequence is young but lively. van Benthem's [21] work on dynamic consequence relations, sketched above, primarily treats changes in *hard information* induced by *public announcements*. To repeat, the dynamic consequence relation that he studies takes the following form:

$$\phi_1,...,\phi_n \models_{dyn} \psi \text{ iff } \mathcal{M},w \models [!\phi_1]...[!\phi_n]\psi \text{ for all pointed models } \mathcal{M},w.$$

He showed that despite violating cut, monotonicity and contraction, this particular type of dynamic consequence relation can be completely axiomatized by a set of rules reminiscent of rules found in the literature on non-monotonic reasoning. This work has set the agenda. Cordón-Franco *et al.* [8] have generalized van Benthem's approach to allow for public announcements in specific informational contexts. To be more precise, their main object of study is a relation defined as follows:

$$\Gamma,\phi_1,...,\phi_n \models_{dyn} \psi \text{ iff } \mathcal{M},w \models [!\phi_1]...[!\phi_n]\psi \text{ for all pointed models } \mathcal{M},w \text{ such that } \mathcal{M},w \models \Gamma.$$

[2] Of course, not all DEL-like systems are prone to this reduction technique. Well-known examples are public announcements in S5 + Common knowledge [5] and the logic of 'epistemic protocols' [24].

[3] For this terminology, see e.g. [10] and the references therein.

This is the 'local' variant of the contextualized dynamic consequence relation. The authors [8] also study the properties of a global version of this relation, which uses model validity instead of truth at pointed models. They have shown that the rules used by van Benthem fail in this more general setting, for both local and global versions of the relation. Following a different trend, Aucher [1, 2] studied dynamic consequences of the following form:

$$\phi_1 ; \phi_2 \vdash \phi_3$$

These should be read 'in every situation where ϕ_1 holds, updating by a ϕ_2-type of learning event will result in a situation where ϕ_3 holds.' Aucher has shown that the set of valid sequents of that form, defined for general dynamic epistemic logic and arbitraty epistemic attitudes, can be completely axiomatized. This is the framework that is adopted in the present article, so we present it in greater detail. The formal definition can be found in Section 4.

The three-place sequents used by Aucher are rigid. They connect pairs of single formulas to single formulas. The left-most formula describes the initial situation in which the update is performed. The second formula describes the type of learning event that is taking place. The formula on the right of the turnstile describes the situation after the update. This framework thus captures a fragment of the set of DEL-validities for a one-shot update. Furthermore, as we shall see below, the two formulas on the left describe different types of structures. The first and the third formulas are interpreted according to standard Kripke models for multi-modal logics. The second one describes so-called 'event models'. Again, we give further details on this terminology below. What matters here is that these two formulas, because they describe different kinds of objects, might be from different languages, so rules like contraction or exchange have no straightforward correspondent in this setup. Some form of weakening can be expressed, but not the standard rule, i.e. adding a formula to the left or the right side.

Both Aucher's and Cordón-Franco et al.'s sequents are contextual dynamic consequence relations. They take as argument formula(s) that describe the initial situation in which the update is performed. In Aucher's framework, this is always a single formula or, equivalently, a finite set thereof, while Cordón-Franco et al. allows for infinite sets of such formulas. As mentioned, the three-place sequents are suited to talking about one-shot updates. van Benthem and Cordón-Franco et al., on the other hand, explicitly address the iterated case. This facilitates the expression and the study of classical structural properties, and both van Benthem and Cordón-Franco et al. do indeed focus on identifying valid and invalid ones. Another important difference is that van Benthem and Cordón-Franco et al. study dynamic consequences for specific update rules: hard and soft public announcements. The consequence relation defined by Aucher builds on full DEL, and thus covers a large variety of information update mechanisms. The set of valid rules identified by Aucher has been proved to axiomatize completely dynamic consequence for DEL.

The present article contributes to this literature by studying dynamic consequences for a large class of soft information update mechanisms. Our underlying dynamic framework is the soft information counterpart of DEL, developed by Baltag and Smets [3]. Its models and language are those of beliefs and conditional beliefs. As such it is more specific than DEL, which works for arbitrary multi-modal systems. The update rule is different, however. While DEL uses the so-called product update rule, Baltag and Smets have proposed a 'lexicographic' update procedure, which gives priority to the information carried by the incoming (soft) epistemic action. They have shown that this update rule embeds all the most well-known update rules for soft information. These are central to the theory of multi-agent belief revision [6, 9, 22]. It is thus important to study the dynamic consequence relations that they generate, which is what we do here. We provide a sound and complete axiomatization for soft information dynamics in the three-place sequents framework just discussed. We then study its generalization to sequences of information update. For that case we focus on identifying which classical structural rules hold and which fail.

3 Lexicographic update for soft information

Here we introduce the framework for studying soft information dynamics developed in [3] or, as the authors put it, a general theory of multi-agent belief revision. We only give the bare essentials, before moving on to dynamic consequence relations. For more details and examples the reader should consult [3] and [22].

3.1 Static models and language

We are interested in the dynamics of so-called *soft* informational attitudes. These are attitudes, like beliefs, which are revisable, might be mistaken, and might not be fully introspective.[4] Our starting point is thus a static modal language expressive enough to encode *conditional beliefs*, i.e. modalities of the form $B_i^{\phi}\psi$, to be read 'conditional on ϕ, agent i believes that ψ'. Mainly for technical reasons, we work in the language $\mathcal{L}_S$ ('S' for 'State model'), in which conditional beliefs are definable. Throughout we assume a finite set A of agents and denote its elements i,j, etc. Let $\mathfrak{P}$ be a countable set of atomic propositions.

$$\phi := p \in \mathfrak{P} \mid \neg\phi \mid \phi\wedge\phi \mid [\sim_i]\phi \mid [\leq_i]\phi$$

We write $\langle\sim_i\rangle\phi$ for $\neg[\sim_i]\neg\phi$ and $\langle\leq_i\rangle\phi$ for $\neg[\leq_i]\neg\phi$. Formulas of the form $[\sim_i]\phi$ and $[\leq_i]\phi$ should be read, respectively, as 'agent i knows that ϕ' and 'agent i safely believes that ϕ'. We write $\langle\sim_i\rangle\phi$ for $\neg[\sim_i]\neg\phi$ and $\langle\leq_i\rangle\phi$ for $\neg[\leq_i]\neg\phi$. Conditional beliefs $B_i^{\phi}\psi$ are then defined as follows:

$$B_i^{\phi}\psi \Leftrightarrow_{df} \langle\sim_i\rangle\phi \rightarrow \langle\sim_i\rangle(\phi\wedge[\leq_i](\phi\rightarrow\psi))$$

This language is interpreted in *plausibility models*, which are Kripke structures [7] equipped with a collection of partial pre-orders $\leq_i$ and a valuation V assigning to each state $w\in W$ a subset of a given set of atomic propositions.

Definition 3.1
Let $\mathfrak{P}$ be a countable set of propositions and A a set of agents. A **plausibility model** $\mathcal{M}$ is a tuple $\langle W, \{\leq_i\}_{i\in A}, V\rangle$ where:

- W is a non-empty set of states.
- for each $i\in A$, $\leq_i$ is a well-founded,[5] reflexive and transitive relation on W.
- $V: W\rightarrow\mathcal{P}(\mathfrak{P})$ is a valuation function.

Plausibility frames are models minus the valuation. The reader familiar with the literature on belief revision or non-monotonic logic will recognize these models as the usual preferential models—see

[4]By 'introspective' we mean attitudes that do not validate the so-called positive and negative instrospection axioms. Positive introspection is the '4' axiom, which states that if an agent believes ϕ, then she believes that she believes ϕ. Negative introspection is the '5' axiom, which states that if an agent does not believe ϕ, then she believes that she does not believe ϕ.
[5]Well-foundedness is only needed to ensure that conditional beliefs are well defined.

e.g. [15] and references therein. This is what they are, with the proviso that this is a multi-agent setting. In this framework, the epistemic accessibility relation usually associated with hard attitudes like knowledge is definable.

Definition 3.2
Let $\mathcal{M}$ be a plausibility model. The **epistemic accessibility relation** $\sim_i$ for agent i is defined as follows:

- $w \sim_i w'$ iff $w \leq_i w'$ or $w' \leq_i w$.

Call the set $[w]_i = \{w' : w' \sim_i w\}$ the similarity class of w.

Definition 3.3
Let $\mathcal{M}$ be a plausibility model. The similarity class $[w]_i$ for a state w is **locally connected** whenever every state $w', w'' \in [w]_i$ are comparable, i.e. either $w' \leq_i w''$ or the other way around. A plausibility model is locally connected if all its similarity classes are.

It is easy to check that $\sim_i$ is an equivalence relation on locally connected models. In this article our attention is confined to such models.

The truth conditions for the modalities then go as follows.

- $\mathcal{M}, w \models [\leq_i]\phi$ iff $\mathcal{M}, w \models \phi$ for all $w' \geq_i w$.
- $\mathcal{M}, w \models [\sim_i]\phi$ iff $\mathcal{M}, w \models \phi$ for all $w' \sim_i w$.

Validity on models and frames, and classes thereof, is defined as usual. Each $[\leq_i]$ is a KT modality. Well-foundedness and transitivity can be enforced by the Löb axiom. Each $[\sim_i]$ is an S5 modality. Interaction between $[\leq_i]$ and $[\sim_i]$ can be captured by standard inclusion axioms. Local connectedness is enforced by the following [3]:

$$[\sim_i](\phi \vee [\leq_i]\psi) \wedge [\sim_i](\psi \vee [\leq_i]\phi) \rightarrow [\sim_i]\phi \vee [\sim_i]\psi \tag{LC}$$

Putting all this together we get:

Theorem 3.4
Multi-agent KTL + S5 together with the relevant inclusion axioms and (LC) is sound and complete with respect to the class of locally connected plausibility frames.

In what follows, we write $\vdash_S \phi$ when ϕ is a theorem of multi-agent KTL + S5 together with the relevant inclusion axioms and (LC).

3.2 Event models and language

The modelling of epistemic action developed in [3] follows the same methodology as for DEL. Soft epistemic actions are encoded in so-called *event models*. These are Kripke structures equipped with a collection of pre-orders, where the elements of its domain are thought of as basic events or programs.

Definition 3.5
Let $\mathfrak{E}$ be a finite set of propositional atoms. An **event model** $\mathcal{E}$ is a tuple $\langle E, \{\leq_i^{\mathcal{E}}\}_{i \in A}, V^{\mathcal{E}} \rangle$ where:

- E is a non-empty, finite set of events.
- for each $i \in A$, $\leq_i^{\mathcal{E}}$ is a well-founded, reflexive, transitive and locally connected relation on E.

- $V^{\mathcal{E}}: E \to Event$ is a valuation function such that $V(e)$ is a singleton for all $e \in E$.

We write $e \cong_i e'$ whenever $e \leq_i e'$ and $e' \leq_i e$. An event frame is an event model minus the valuation.

This presentation, taken from [1], differs from the usual one. In DEL, event models are equipped with a precondition function instead of a propositional valuation. This function is meant to specify the conditions of executability of events or epistemic actions. However, to define the dynamic consequence relation we need a language to describe event models and types of epistemic actions. To do that it is more convenient to work with a propositional valuation. This valuation is constrained to re-capture the precondition function.

DEFINITION 3.6
Let $\mathcal{E}$ be a finite set of atomic propositions. A **precondition function** is a function $pre: \mathcal{E} \to \mathcal{L}_S$ that assigns to each $e \in \mathcal{E}$ a formula of $\mathcal{L}_S$.

From now on, when we talk about an event model $\mathcal{E}$ we will implicitly assume a precondition function on a given finite set of atomic propositions $\mathcal{E}$ associated with $\mathcal{E}$.

As just mentioned, the dynamic consequence relation requires a language to talk about event models. Unlike the language for state models, this one is equipped with a Kripke modality for the strict sub-relation of the plausibility order $\leq_i$. The reason for that is the lexicographic update rule. It makes explicit reference to some events being strictly more plausible than others, and this is reflected in the set of valid sequents that we want to capture.

The language $\mathcal{L}_{\mathcal{E}}$ is constructed as follows. To simplify the notation we also use $e, e', \ldots$ to refer to elements of the set $\mathcal{E}$ of atomic propositions.

$$\phi := e \in \mathcal{E} \mid \neg\phi \mid \phi \wedge \phi \mid [\sim_i]^{\mathcal{E}}\phi \mid [\leq_i]^{\mathcal{E}}\phi \mid [<_i]^{\mathcal{E}}\phi$$

We write $\langle\sim_i\rangle^{\mathcal{E}}\phi$ for $\neg[\sim_i]^{\mathcal{E}}\neg\phi$, $\langle\leq_i\rangle^{\mathcal{E}}\phi$ for $\neg[\leq_i]^{\mathcal{E}}\neg\phi$ and $\langle<_i\rangle^{\mathcal{E}}\phi$ for $\neg\langle<_i\rangle^{\mathcal{E}}\neg\phi$. Truth for this language is defined as usual, and so are the different notions of validity.

DEFINITION 3.7
Let $\mathcal{E}, e$ be an event model.

- $\mathcal{E}, e \models [\leq_i]\phi$ iff $\mathcal{E}, e' \models \phi$ for all $e' \geq_i e$.
- $\mathcal{E}, e \models [<_i]\phi$ iff $\mathcal{E}, e' \models \phi$ for all $e' >_i$ e.
- $\mathcal{E}, e \models [\sim_i]\phi$ iff $\mathcal{E}, e' \models \phi$ for all $e' \sim_i e$.

The validities for the weak and strict boxes $[\leq_i]$ and $[<_i]$ and their proper interrelation over the class of event frames is completely axiomatized by the system $\Lambda^{\mathcal{L}_{\mathcal{P}}}$ presented in [25, p.93], again together with the Löb axiom to take care of well-foundedness. The interaction between these modalities and the S5 box $[\sim_i]$ is a simple inclusion. It can easily be dealt with. We assume it when we write $\Lambda^{\mathcal{L}_{\mathcal{P}}}$+S5. The constraint on the valuation function is captured by a set $Ex_{\mathfrak{E}}$ of axioms of the following form, one for each $e \in \mathfrak{E}$.

$$\vdash e \leftrightarrow \neg \bigvee \mathfrak{E}_{-e} \tag{Excl}$$

where $\mathfrak{E}_{-e}$ is the set of atoms in $\mathfrak{E}$ except e.

THEOREM 3.8
Multi-agent $\Lambda^{\mathcal{L}_{\mathcal{P}}}$+S5+$Ex_{\mathfrak{E}}$+(LC) is sound and complete with respect to the class of event frames.

In what follows, we write $\vdash_E \phi$ when ϕ is a theorem of multi-agent $\Lambda^{\mathcal{L}_{\mathcal{P}}}$+S5+$Ex_{\mathfrak{E}}$.

3.3 Lexicographic update

The lexicographic update rule takes pairs of plausibility and event models $\mathcal{M},w$ and $\mathcal{E},e$ and return the updated model $\mathcal{M}\otimes\mathcal{E}$ where the domain is the set of pairs (w,e) such that $\mathcal{M},w$ satisfies the precondition of e, written $\mathcal{M},w\models pre(e)$ and the valuation is taken directly from $\mathcal{M}$, i.e $V'(w,e)=V(w)$. The adjective 'lexicographic' comes from the update rule for the pre-orders $\leq_i$, which gives priority to the events:

DEFINITION 3.9
Let $\mathcal{E}$ be an event model and *pre* a precondition function on $\mathfrak{E}$. The **lexicographic update** $\mathcal{M}\otimes\mathcal{E}=\langle W',\leq_i',V'\rangle$ is defined as follows:

- $W'=\{(w,e)|\mathcal{M},w\models pre(e)$ for the unique e such that $\mathcal{E},e\models e\}$.
- $(w,e)\leq_i'(w',e')$ iff either $e<e'$ and $w\sim_i w'$ or $e\cong_i e'$ and $w\leq w'$.
- $V'(w,e)=V(w)$.

4 Dynamic consequence for lexicographic update

In this section, we build up to one of the two main contributions of this article: a sound and complete system of rules for the soft dynamic consequence relation. We first give a semantic definition of this consequence relation. Then we introduce a set of axioms and rules, and show that they are sound. The proof that they are also complete uses the same technique as in [1], namely a detour via 'canonical formulas' for plausibility and event models. We introduce and prove basic facts about these formulas first, and then progress to the completeness result.

4.1 Semantic dynamic consequence relation

Dynamic consequence relations for lexicographic update are defined as follows:

DEFINITION 4.1 (Soft Dynamic Consequence Relation - Semantic Version)
Let ϕ_1 and ϕ_3 be formulas of $\mathcal{L}_S$ and ϕ_2 a formula of $\mathcal{L}_E$. A *soft dynamic sequent* is an object of the following form

$$\phi_1;\phi_2\models\phi_3$$

We say that such a sequent is valid whenever for all M,w and E,e, if $M,w\models\phi_1\wedge pre(e)$ and $E,e\models\phi_2$ then $(\mathcal{M}\otimes\mathcal{E},(w,e))\models\phi_3$

So we are working with three-place sequents in the same fashion as in [1]. We emphasize again that the second formula on the left, ϕ_2, is not of the same language as the two other formulas ϕ_1 and ϕ_3. In other words, in this framework there is no straightforward version of exchange between ϕ_1 and ϕ_2, or contraction for that matter. Some form of cut can be formulated here, but we shall return to this in Section 5.

This dynamic consequence relation is contextualized in the same sense as those studied by [8]. The rule takes as first argument a formula, ϕ_1, which describes the initial situation in which the update takes place. This could of course be equivalently formulated by taking a finite set of formulas in the first argument set. We do that in Section 5, when the connection with substructural systems will be made more explicit. Unlike [8], however, we do not consider initial situations described by infinite sets of formulas.

This dynamic consequence relation takes as second argument general types of epistemic actions or learning events. A sequent of the form $\phi_1;\phi_2\models\phi_3$ is valid iff *any* update of type ϕ_2 in an initial

TABLE 1. Axioms and rules for the dynamic consequence relation $\vdash$. In R3–R5 it must be the case that $\vdash_L \psi_i \to \phi_i$ for either L=S and $i = 1,3$ or L=E and $i = 2$.

$\bot;\phi_2 \vdash \phi_3$	A1	$\phi_1;\bot \vdash \phi_3$	A2
$\phi_1;\phi_2 \vdash \top$	A3	$p;\phi_2 \vdash p$	A4
$\neg p;\phi_2 \vdash \neg p$	A5	$\neg Pre(e);e \vdash \bot$	A6

$$\frac{\phi_1;\phi_2 \vdash \phi_3 \quad \phi_1;\phi_2 \vdash \phi_4}{\phi_1;\phi_2 \vdash \phi_3 \wedge \phi_4}\ R1$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3 \quad \psi_1;\phi_2 \vdash \phi_4}{\phi_1 \vee \psi_1;\phi_2 \vdash \phi_3}\ R2 \qquad \frac{\phi_1;\phi_2 \vdash \phi_3 \quad \phi_1;\psi_2 \vdash \phi_4}{\phi_1;\phi_2 \vee \psi_2 \vdash \phi_3}\ R2'$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{\psi_1;\phi_2 \vdash \phi_3}\ R3 \quad \frac{\phi_1;\phi_2 \vdash \phi_3}{\phi_1;\psi_2 \vdash \phi_3}\ R4 \quad \frac{\phi_1;\phi_2 \vdash \psi_3}{\phi_1;\phi_2 \vdash \phi_3}\ R5$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3 \quad \phi_4;\phi_5 \vdash \phi_3}{[\sim_i]\phi_1 \wedge [\leq_i]\phi_4;[<_i]^{\mathcal{E}}\phi_2 \wedge [\leq_i]^{\mathcal{E}}\phi_5 \vdash [\leq_i]\phi_3}\ (R6)$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{\langle\sim_i\rangle(\phi_1 \wedge pre(e));\langle<_i\rangle^{\mathcal{E}}(\phi_2 \wedge e) \vdash \langle\leq_i\rangle\phi_3}\ (R7)$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{\langle\leq_i\rangle(\phi_1 \wedge pre(e));\langle\leq_i\rangle^{\mathcal{E}}(\phi_2 \wedge e) \vdash \langle\leq_i\rangle\phi_3}\ (R8)$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{[\sim_i]\phi_1;[\sim_i]^{\mathcal{E}}\phi_2 \vdash [\sim_i]\phi_3}\ (R9)$$

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{\langle\sim_i\rangle(\phi_1 \wedge pre(e));\langle\sim_i\rangle^{\mathcal{E}}(\phi_2 \wedge e) \vdash \langle\sim_i\rangle\phi_3}\ (R10)$$

situation where ϕ_1 holds results in a ϕ_3-situation. So our object of study here is inferences valid for large collections of soft information updates, as opposed to inferences holding for specific epistemic actions, such as public announcements or conservative upgrades.

4.2 The Gentzen-style system

The set of soft dynamic sequents is completly axiomatized by the set of axioms and rules presented in Table 1. Most of them are directly imported from [1]. At the level of axioms, for instance, the set of validities is the same for DEL and its soft version, which we consider in this article. Epistemic actions do not change the truth value of atoms in either framework (A4, A5), and a pointed event model is executable in a pointed plausibility model whenever the latter satisfies the preconditions of the former (A6). Soft dynamic consequence also inherits from DEL rules to introduce conjunction to the right (R1), and disjunction to the left (R2, R2').

The rules R3 to R5 are also ported directly from [1], but a few remarks are in order. Each of these three rules makes explicit reference to valid implications in the "underlying logic". Take for instance R3:

$$\frac{\phi_1;\phi_2 \vdash \phi_3}{\psi_1;\phi_2 \vdash \phi_3}\ R3$$

This rule holds with the proviso that $\vdash_S \psi_1 \to \phi_1$. This might sound suspicious to the proof-theoretically inclined reader. Are we not trying to capture the set of soft-DEL valid formulas? No, not all of them. To see why, one must bear in mind that we are axiomatizing here the set of valid soft dynamic sequents. These, in turn, correspond to a fragment of the dynamic extension of the language for safe belief and knowledge. Take a soft dynamic sequent $\phi_1;\phi_2 \models \phi_3$. Let $\mathcal{E},e$ be any pointed event model such that $\mathcal{E},e \models \phi_2$. The sequent is valid iff for all pointed models $\mathcal{M},w$, we have that

$$\mathcal{M},w \models \phi_1 \to [\mathcal{E},e]\phi_3$$

So axiomatizing soft dynamic consequence is axiomatizing a fragment of the validities in $\mathcal{L}_S$. The key operator of this fragment is the one for lexicographic update. Its behaviour presupposes the complete logic of the respective modal operators for static and event models. So using valid implications from these underlying logics in the Gentzen-style system is not circular. The system builds on these logics, as opposed to trying to capture them.

Rules $R3$–$R5$ are also important because they are as close to Weakening as one can get in the three-place sequents we are working with. The classical form of weakening allows one to add formulas on the left, and also on the right if one allows for multiple conclusions. We cannot do this in our three-place sequents. What we can do is replace existing formulas with logically weaker or stronger ones, depending on whether they occur on the right or on the left side of the sequent. Take R3 again. It states that if one can prove the sequent $\phi_1;\phi_2 \vdash \phi_3$, then one can prove the overall logically weaker sequent $\psi_1;\phi_2 \vdash \phi_3$ by replacing ϕ_1 in the antecedent with any logically stronger formula ψ_1. By logically stronger we mean that the latter implies the former in the underlying logic.

Of course, the specificity of our axiomatization lies in the rules R6–R10, which reflects the behaviour of the lexicographic update rule. The two-premise rule R6 introduces boxes in front of the different arguments in the sequent. Dually, the pair R7 and R8 allows the introduction of diamonds. This reflects the disjunctive character of lexicographic update. Take the diamond-introduction rules, for instance. Informally they can be read as follows. Suppose any update by a ϕ_2-type of epistemic action in a situation where ϕ_1 holds will result in a situation where ϕ_3. When will a ϕ_3-state (w',e) be $\leq_i$-reachable from another state (w,e) after an update? By assumption and the lexicographic update rule we know that this can happen in two cases. The first one is when there was a ϕ_1 state w' comparable to w and a ϕ_2-event e' strictly more plausible than e, under the assumption that e' was executable in w'. This is what $R7$ states. The second one is when e' was a ϕ_2-state equi-plausible to e and w' a ϕ_1 at least as plausible as w, with the same proviso regarding preconditions. $R8$ covers this case because equi-plausibility implies comparability. See the proof of soundness below for details. The same idea works dually for $R6$. The reader can check that, because of the way lexicographic update is set up, it collapses to the standard DEL product update for the hard information modality $[\sim_i]$. It is therefore not surprising that the rules regimenting the introduction of that modality to the right are exactly as in [1]. With these informal explanations in mind, we now proceed to show soundness of these rules.

Observation 4.1
All axioms and rules in Table 1 are sound.

Proof. We only prove the cases for R6–R10. Soundness for the axioms and for R1–R5 proceeds exactly as in [1].

R6 Suppose $\mathcal{M},w \models [\sim_i]\phi_1 \wedge [\leq_i]\phi_4$ and $\mathcal{E},e \models [<_i]^{\mathcal{E}}\phi_2 \wedge [\leq_i]^{\mathcal{E}}\phi_4$. It is sufficient to show that $\mathcal{M}\otimes\mathcal{E},(w,e) \models [\leq_i]\phi_3$. If the $\mathcal{M},w \not\models pre(e)$ then we are done. Suppose then that $\mathcal{M},w \models$

pre(e). Take any $(w_1,e_1)\in\mathcal{M}\otimes\mathcal{E}$ such that $(w_1,e_1)\geq_i(w,e)$. This can happen in two cases. First, it can be that $w_1\sim_i w$ and $e_1>_i e$. Because $\mathcal{M},w\models[\sim_i]\phi_1$ and $\mathcal{E},e\models[<_i]^{\mathcal{E}}\phi_2$, respectively, we get $\mathcal{M},w_1\models\phi_1$ and $\mathcal{E},e_1\models\phi_2$. But since $\phi_1;\phi_2\models\phi_3$, we get that $\mathcal{M}\otimes\mathcal{E},(w_1,e_1)\models\phi_3$. The other case is when $e_1\cong_i e$ and $w_1\leq_i w$. Since $\mathcal{M},w\models[\leq_i]\phi_4$ and $\mathcal{E},e\models[\leq_i]^{\mathcal{E}}\phi_5$ we get $\mathcal{M},w_1\models\phi_4$ and $\mathcal{E},e_1\models\phi_5$. But then since $\phi_4;\phi_5\models\phi_3$, we get $\mathcal{M}\otimes\mathcal{E},(w_1,e_1)\models\phi_3$ again. So we have just shown that the latter holds for any $(w_1,e_1)\geq_i(w,e)$, which means $\mathcal{M}\otimes\mathcal{E},(w,e)\models[\leq_i]\phi_3$.

R7 (Sketch) $\mathcal{E},e\models\langle<_i\rangle^{\mathcal{E}}(\phi_2\wedge e')$ gives us the e' where ϕ_2 (and e' hold) such that $e<_i e'$ and similarly for the required w' from $\mathcal{M},w\models\langle\sim_i\rangle(\phi_1\wedge pre(e'))$. So the pair (w',e') turn out to be related to (w,e) in the updated model. But then we can use $\phi_1;\phi_2\models\phi_3$ to conclude that $\mathcal{M}\otimes\mathcal{E},(w',e')\models\phi_3$.

R8 (Sketch) $\mathcal{E},e\models\langle\leq_i\rangle^{\mathcal{E}}(\phi_2\wedge e')$ gives us the e' where ϕ_2 and e' hold such that $e\leq_i e'$, and $\mathcal{M},w\models\langle\leq_i\rangle(\phi_1\wedge pre(e'))$ the $w'_i\geq w$ that satisfy ϕ_1 and would survive the update with e'. Now, if the relation between e and e' does not go the other way around, the argument follows the same line as for R7. If it does, then (w',e') are related to (w,e) in the updated model, from which we can conclude $\mathcal{M}\otimes\mathcal{E},(w',e')\models\phi_3$ from $\phi_1;\phi_2\models\phi_3$.

R9 Suppose $\mathcal{M},w\models[\sim_i]\phi_1$ and $\mathcal{E},e\models[\sim_i]^{\mathcal{E}}\phi_2$. Take any $(w_1,e_1)\in\mathcal{M}\otimes\mathcal{E}$ such that $(w,e)\sim_i(w_1,e_1)$, i.e. either $(w,e)\leq_i(w_1,e_1)$ or $(w_1,e_1)\leq_i(w,e)$. Consider the first case. There are two sub-cases. First $e<_i e_1$ and $w\sim_i w_1$. From the latter we get directly that $\mathcal{M},w\models\phi_1$. From the latter it follows that $e\sim_i e_1$, and so $\mathcal{M},e\models\phi_2$. But since $\phi_1;\phi_2\models\phi_3$, we get that $\mathcal{M}\otimes\mathcal{E},(w_1,e_1)\models\phi_3$. The second sub-case is when $e\cong_i e_1$ and $w\leq_i w_1$. Each entails its respective weaker $\sim_i$, from which the desired conclusion follows. Now back up one level. The second case is when $(w_1,e_1)\leq_i(w,e)$. Again, there are two sub-cases. They follow exactly the same steps because the relation $\sim_i$ abstracts away from the direction of $\leq_i$.

R10 Take any $\mathcal{M},w$ and $\mathcal{E},e$ such that $\mathcal{M},w\models pre(e)$, $\mathcal{M},w\models\langle\sim_i\rangle(\phi_1\wedge pre(e'))$ and $\mathcal{E},e\models\langle\sim_i\rangle^{\mathcal{E}}(\phi_2\wedge e')$. We have to show that $\mathcal{M}\otimes\mathcal{E},(w,e)\models\langle\sim_i\rangle\phi_3$. Since $\mathcal{M},w\models pre(e)$ we know that the pair (w,e) is in the updated model. We know furthermore that there is a w' such that $w\sim_i w'$ and $\mathcal{M},w'\models\phi\wedge pre(e')$. So (w',e') is in the updated model too. Do we have $(w,e)\sim_i(w',e')$? There are a number of cases to consider. We can group them in two families. First, whenever $e<_i e'$ or the other way around we get the desired $(w,e)\sim_i(w',e')$. In that case, because $w\sim_i w'$ the update rule gives priority to events. The remaining group of cases is when $e\cong_i e'$. But then the update rule reverts to the order of the relation in the state model, but whatever that is we get $(w,e)\sim_i(w',e')$. But then we are done, for since $\phi_1;\phi_2\models\phi_3$ we know that $\mathcal{M}\otimes\mathcal{E},(w',e')\models\phi_3$ and so $\mathcal{M}\otimes\mathcal{E},(w,e)\models\langle\sim_i\rangle\phi_3$.

■

4.3 Completeness for the Gentzen-style system

In this section, we prove that the set of axioms and rules in Table 1 completely axiomatize the set of valid soft dynamic sequents. First we define formally what a derivable soft dynamic sequent is.

DEFINITION 4.2
A soft dynamic sequent is *derivable*, written $\phi_1,\phi_2\vdash\phi_3$, iff it is either an instance of one of the axioms A1-A7 or it follows from such an axiom by a finite number of applications of the rules R1–R10.

As in [1], the completeness proof proceeds in two steps. First we prove that all valid sequents formed of so-called 'canonical formulas' are derivable. Then we show that this result extends to arbitrary soft dynamic sequents. So we need to explain first what these canonical formulas are.

4.3.1 Canonical formulas and cover modality

Canonical modal formulas have been known at least since Kit Fine's 1975 paper. The formulation we take here comes from Larry Moss [17]. We present it independently of the specific modal operators in $\mathcal{L}_S$ and $\mathcal{L}_E$. The facts we mention about them hold in general. We thus speak abstractly of a language $\mathcal{L}$ and of boxes $[R_i]$ and diamonds $\langle R_i\rangle$ for $i \in A$ and $R \in \{\sim_i, \leq_i, \sim_i^{\mathcal{E}}, \leq_i^{\mathcal{E}}, <_i^{E}\}$. Similarly, when we use theorems of the logic of plausibility or event models, we omit the subscripts on $\vdash_S$ and $\vdash_E$ unless ambiguity can arise. We only present proofs that are in any interesting way different from those in [17].

Definition 4.2
Let P be a finite subset of $\mathfrak{P}$ and π a subset of P. Let Φ, Φ_i^R be finite subsets of $\mathcal{L}$, with R in either $\{\sim_i, \leq_i,\}_{i\in A}$ or $\{\sim_i^{\mathcal{E}}, \leq_i^{\mathcal{E}}, <_i^{E}\}_{i\in A}$, depending on the models at hand. Then:

$$\odot\pi := \bigwedge_{p\in\pi} p \wedge \bigwedge_{p\in P-\pi} \neg p$$

$$\nabla_{R_i}\Phi := [R_i] \bigvee_{\phi\in\Phi} \phi \wedge \bigwedge_{\phi\in\Phi} \langle R\rangle\phi$$

$$\pi \bullet \bigwedge_{R_i} \Phi_i^R := \odot\pi \wedge \bigwedge_{R_i} (\nabla_{R_i}\Phi_i^R)$$

The 'bullet' ($\bullet$) and 'nabla' (∇) notations are from [28]. We use them for convenience.

Canonical formulas are built inductively. The intuition behind the construction is to view canonical formulas as descriptions of pointed models. Canonical formulas of level 0 describe the propositional valuation up to the given finite subset P. Canonical formulas of level $n \geq 1$ list all and only the valuations accessible in n step from a given state/valuation. The construction is essentially the same for epistemic-plausibility and for event models. The only difference is for the base case.

Definition 4.3 (Canonical formula of $\mathcal{L}_S$)
Let P be a non-empty finite subset of $\mathfrak{P}$.

(1) $S_P^0 = \{\odot\pi \mid \pi \subseteq P\}$.
(2) $S_P^{n+1} = \{\pi \bullet \bigwedge_{R_i} \Phi_i^R \mid \pi \subseteq P, \Phi_i^R \subseteq S_P^n\}$.

Call any formula $\delta_n \in S_P^n$ a *canonical* $\mathcal{L}_S$*-P-formula of depth n*.

Definition 4.4 (Canonical formula of $\mathcal{L}_{\mathcal{E}}$)
Recall that $\mathfrak{E}$ is already a finite set, so we start with that.

(1) $S_P^0 = \{\odot\pi \mid \pi$ is a singleton subset of $\mathfrak{E}\}$.
(2) $S_P^{n+1} = \{\pi \bullet \bigwedge_{R_i} \Phi_i^R \mid \pi \subseteq P, \Phi_i^R \subseteq S_P^n\}$.

Call any formula $\delta_n \in S_P^n$ a *canonical* $\mathcal{L}_{\mathcal{E}}$*-P-formula of depth n*.

Except when ambiguity can arise we drop the prefix '$\mathcal{L}_{S/\mathcal{E}}$-P' and just say 'canonical formulas.'

LEMMA 4.5
Let M,w be an epistemic-plausibility model and P a finite set of atomic propositions. Then for all n there is exactly one $\delta \in S_P^n$ such that $M,w \models \delta$. The same holds, *mutatis mutandis*, for event models.

PROOF. Essentially the same as in [17], except for a minor detail in the case of event models. ■

Now we can state the crucial fact about canonical formulas. The argument only uses normality of the modal operator at hand. So we omit it.

LEMMA 4.6
Let ϕ be a formula of $\mathcal{L}_S$ of modal depth n, and let $P(\phi) = \{p \in \mathfrak{P} : p \text{ is a sub-formula of } \phi\}$.

- For all $\delta \in S^n_{P(\phi)}$, either $\vdash \delta \to \phi$ or $\vdash \delta \to \neg\phi$.
- $\vdash_S \bigvee_{S^n_{P(\phi)}} \delta$

The same holds, *mutatis mutandis*, for event models.

The main point of the previous lemma is to get us the following corollary, which will be crucial in lifting the completeness result for canonical formulas to arbitrary soft dynamic sequents. We omit the proof, as it is essentially the same as that in [1]. The 'S possibly empty' proviso is to take care of the case where ϕ is not satisfiable, following the usual convention that $\bigvee \emptyset = \bot$.

COROLLARY 4.7
Let ϕ be a formula of $\mathcal{L}_S$ of modal depth n, and let $P(\phi) = \{p : p \text{ is a sub-formula of } \phi\}$. Then there is a $S \subseteq S^n_{P(\phi)}$, with S possibly empty, such that :

$$\vdash \phi \leftrightarrow \bigvee_{\delta \in S} \delta$$

4.3.2 Canonical formulas and n-bisimilarity

Bisimulation is a core concept for the model theory of modal logic, cf. [7]. So-called n-bisimulation is a finite version of that notion. Canonical formulas are closely connected to n-bisimilarity. This will come in handy below, so we present them in detail.

DEFINITION 4.8 (n-P-Bisimulation)
Let $P \subseteq \mathfrak{P}$. Two plausibility models $\mathcal{M}, w$ and $\mathcal{M}', v$ are *n-P-bisimular*, written $(\mathcal{M}, w) \underline{\leftrightarrow}_n (\mathcal{M}', w')$ whenever there is a sequence $\underline{\leftrightarrow}_n \subseteq \ldots \subseteq \underline{\leftrightarrow}_0$ of relations $\underline{\leftrightarrow}_i \subseteq W \times W'$ such that, for all $k+1 \leq n$, $w' \in W$, $v' \in W'$ and $R_i \in \{\leq_i, \sim_i\}$

(1) $w \underline{\leftrightarrow}_n v$;
(2) If $w' \underline{\leftrightarrow}_0 v'$ then $(V(w') \cap P) = (V'(v') \cap P)$.
(3) If $w' \underline{\leftrightarrow}_{k+1} v'$ if $w' R_i w''$ then there is a $v'' \in W'$ such that $v' R'_i v''$ and $w'' \underline{\leftrightarrow}_k v''$.
(4) Same from W' to W.

This definition is adapted to event models in the obvious way. Now suppose that two plausibility or event models satisfy the same canonical formula of degree n. We show now that this is both necessary and sufficient to guarantee that these models are n-bisimilar. It is well-known that bisimulation and n-bisimulation are equivalence relations between models. Canonical formulas capture this equivalence syntactically.

FACT 4.9
Let $P \subseteq_{fin} \mathfrak{P}$. Then t.f.a.e:

(1) There is a $\delta \in S_n^P$ such that $\mathcal{M}, w \models \delta$ and $\mathcal{M}', w' \models \delta$.
(2) $\mathcal{M}, w \underline{\leftrightarrow}_n^P \mathcal{M}', w'$

And similarly for event models.

PROOF. The direction from 2 to 1 follows from the well-known invariance of modal formulas of bounded depths under n-bisimulation and the Lemma 4.5 that each state in a model satisfies exactly one formula in S_n^P. For the other direction take the $\delta \in S_P^n$ that is true in both pointed models. We define the n-P-bisimulation as follows. First set $\underline{\leftrightarrow}_n = \{(w, w')\}$. Then, for $(n > k \geq 0)$, define $\underline{\leftrightarrow}_k$ as $v \underline{\leftrightarrow}_k v'$ iff there is a $\delta' \in S_k^P$ such that $\mathcal{M}, v \models \delta'$ and $\mathcal{M}', v' \models \delta'$. Now we have to check that this is indeed an n-bisimulation between w and w'. We get trivially $w \underline{\leftrightarrow}_n w'$. The case of $w \underline{\leftrightarrow}_0 v'$ follows from the definition of canonical formulas. Now suppose that $v \underline{\leftrightarrow}_{k+1} v'$ for $k > 0$ and there is a v_1 such that $v \leq_i v_1$. Take $\delta \in S_i^{k+1}$ such that $\mathcal{M}, v \models \delta$. By definition of δ we know that there is (exactly) one $\delta' \in S_i^k$ such that $\mathcal{M}, v \models \langle \leq_i \rangle \delta'$ and $\mathcal{M}, v_1 \models \delta'$. But by definition of δ we know that $\mathcal{M}', v' \models \langle \leq_i \rangle \delta'$ too. But then there is a $v' \leq_i v_1'$ such that $\mathcal{M}', v' \models \delta'$, which means by definition that $\mathcal{M}, v_1 \underline{\leftrightarrow} \mathcal{M}', v'$. The argument is the same for the other relations, for the other direction, and for event models. ■

Another important fact about n-bisimulations is that they are preserved by lexicographic update. Take two n-bisimilar pointed event models that can be executed in two n-bisimilar pointed plausibility models. The results of their respective lexicographic update will also be n-bisimilar.

FACT 4.10
Take two state models such that $(\mathcal{M}, w) \underline{\leftrightarrow}_{n+N} (\mathcal{M}', w')$ and event models executable in $\mathcal{M}, w$ such that $(\mathcal{E}, e) \underline{\leftrightarrow}_n (\mathcal{E}', e')$ with $N = max\{md(pre(e)) : e \in E\}$. Then $(\mathcal{M} \otimes \mathcal{E}, (w, e)) \underline{\leftrightarrow}_n (\mathcal{M}' \otimes \mathcal{E}', (w', e'))$.

PROOF. We define the n-bisimulation component-wise:

$$\text{for } 0 \leq k \leq n \text{ let } (w, e) \underline{\leftrightarrow}_k (w', e') \text{ iff } w \underline{\leftrightarrow}_{k+N} w' \text{ and } e \underline{\leftrightarrow}_i e'.$$

Now we have to show that this is an n-bisimulation. Since valuation is preserved in the updated model, we know that $\underline{\leftrightarrow}_0$ preserves propositional equivalence. Now take two $(w, e), (w', e')$ that are related k-bisimilar for $0 < k \leq n$, and suppose that there is a $(w, e) \leq_i (v, f)$. This can happen in two cases.

First, $w \sim_i v$ and $f <_i e$. Since $w \underline{\leftrightarrow}_{k+N} w'$ we know that there is a $v' \sim_i w'$ such that $v' \underline{\leftrightarrow}_{(i+K)-1} v$. Since the pair v, f is in the updated model we know furthermore that $\mathcal{M}, v \models pre(f)$. But observe that since $k > 0$, by our choice of N we can conclude that $\mathcal{M}', v' \models pre(f)$ from $v' \underline{\leftrightarrow}_{(k+N)-1} v$. On the other hand, since $e \underline{\leftrightarrow}_k e'$ we know that there is an $e' < f'$ such that $f' \underline{\leftrightarrow}_{k-1} f$. Since $\mathcal{M}', v' \models pre(f)$ the pair (v', f') will be in the updated model. But now $e' < f'$ and $w' \sim_i v'$ give us together that $(w', e') \leq_i (v', f')$, and $v' \underline{\leftrightarrow}_{(k+N)-1} v$ with $f' \underline{\leftrightarrow}_{k-1} f$ give $(v, f) \underline{\leftrightarrow}_{k-1} (v', f')$.

The second case is when $e \cong_i f$ and $w \leq_i v$. As before, we have to find the right pair (v', f') that is k-bisimilar to (v, f) and $\leq_i$-related to (w', e'). We get to v' in the same way as in the previous case. Now for f'. Since $e \cong_i f$ we know that $e \leq_i f$. Because $e \underline{\leftrightarrow}_k e'$ there must be a f' such that $e' \leq_i f'$ and $e \underline{\leftrightarrow}_{k-1} e'$. From there the argument proceeds as before, except when it comes to establishing that $(w', e') \leq_i (v', f')$. Here we make a case distinction. Suppose we also have $f' \leq_i e'$, then the second case of the lexicographic update rule applies and we conclude that $(w', e') \leq_i (v', f')$. If not, then what we have is $e' <_i f'$. But then since $w' \leq_i v'$ we are back in the first case and can conclude that $e' \leq_i f'$ as well.

■

4.3.3 Proof of the completeness theorem

We are almost ready to start the completeness argument itself. We only need the following preparatory Lemma, which serves as a bridge in the completeness proof. The Lemma essentially states that the failure of a dynamic consequence $\delta_1;\delta_2 \models \phi$ with canonical formulas in the antecendent can be strengthened to a valid dynamic sequent that derives the negation of ϕ.

LEMMA 4.11
Let P' be a subset of $\mathfrak{E}$ and $\delta_2 \in S_n^{P'}$. Take furthermore $\phi \in \mathcal{L}_S$ with $md(\phi) \leq n$, and set $P = P(\phi) \cup \{P(pre(e)) : e \in P'\}$. Take $N = max\{md(pre(e)) : e \in P'\}$, $\delta_1 \in S_P^{n+N}$, and write $pre(\delta_2)$ for $pre(e')$ with e' such that $\vdash \delta_2 \rightarrow e'$.[6] Then t.f.a.e.:

- $\delta_1;\delta_2 \not\models \phi$
- $\delta_1;\delta_2 \models \neg\phi$ and $\models \delta_1 \rightarrow pre(\delta_2)$

PROOF. The direction from (2) to (1) is just unpacking the definitions. For the converse, suppose there is a $\mathcal{M}, w \models \delta_1$ and $\mathcal{E}, e \models \delta_2$ such that $\mathcal{M} \otimes \mathcal{E}, (w,e) \models \neg\phi$. Take any other models $\mathcal{M}', w' \models \delta_1$ and $\mathcal{E}', e' \models \delta_2$. By Lemma 4.9 we know that they are, respectively, $n+N$- and n-bisimilar to $\mathcal{M}, w$ and $\mathcal{E}, e$. But then by Fact 4.10 and Lemma 4.9 again we get that $\mathcal{M}' \otimes \mathcal{E}', (w', e') \models \neg\phi$. $\models_S \delta_1 \rightarrow pre(\delta_2)$ follows from our choice of δ_1 and Lemma 4.6. ■

We are now ready to start the completeness theorem. As mentioned, we prove it first for the canonical formula, and then lift the result to arbitrary soft dynamic sequents.

THEOREM 4.12 (Completeness for canonical formulas)
Let δ_1 and δ_2 be as in Lemma 4.11. Take $\delta_3 \in S_P^n$. Then:

$$\text{If } \delta_1;\delta_2 \models \delta_3 \text{ then } \delta_1;\delta_2 \vdash \delta_3$$

PROOF. By induction on n. The basic case is the same as in [1]. Induction step. We prove the contrapositive. Suppose that $\delta_1;\delta_2 \nvdash \delta_3$. By the shape of δ_3 and Rule R1 one of the following must be the case:

$$\delta_1;\delta_2 \nvdash \odot\pi \tag{4.1}$$

$$\delta_1;\delta_2 \nvdash \bigwedge_{R_i}(\nabla_{R_i}\Phi_i^R) \tag{4.2}$$

Only the second case differs from the argument in [1]. By the shape of $\nabla_{R_i}\Phi_i^R$, using R1 again, one of the following must hold for some R_i and $\Phi_i^R \subseteq S_P^{n-1}$:

(1) $\delta_1;\delta_2 \nvdash \langle R_i\rangle\delta_3'$ for some $\delta_3' \in \Phi_i^R$
(2) $\delta_1;\delta_2 \nvdash [R_i]\bigvee\Phi_i^R$

So there are two cases to consider, depending on whether the failure of the dynamic consequence comes from $\leq_i$ or $\sim_i$. Each has two sub-cases. We consider them in turn.

---

[6] Recall that by construction there will always be a unique such e'.

(1) $R_i = \sim_i$. Both sub-cases here are exactly as in Aucher's [2] paper.

(2a.) $R_i = \leq_i$ $\delta_1;\delta_2 \not\vdash \langle \leq_i \rangle \delta'_3$ for some $\delta'_3 \in \Phi_i^{\leq_i}$. Take such a δ'_3. We have two (similar) intermediate goals. First we want to show that $\delta'_1;\delta'_2 \models \neg\delta'_3$ for any $\delta'_1 \in \Phi_{\delta_1}^{\sim_i} \subseteq S_P^{n-1}$ and $\delta'_2 \in \Phi_{\delta_2}^{<_i} \subseteq S_P^{n-1}$. Take such a δ'_2. By construction we have that $\vdash_{\mathcal{E}} \delta'_2 \to e$ for a unique e. Now take a $\delta'_1 \in \Phi_{\delta_1}^{\sim_i}$. If $\vdash_S \delta'_1 \to \neg pre(e)$ then $\delta'_1;\delta'_2 \models \neg\delta_3$ follows by the same argument as in Aucher's paper. So now suppose that $\vdash_S \delta'_1 \to pre(e)$. By our choice of δ'_1 and δ'_2, applying R3 and R4 twice, we first get that that $\langle \sim_i \rangle \delta'_1; \langle <_i \rangle \delta'_2 \not\vdash \langle \leq_i \rangle \delta'_3$ and then $\langle \sim_i \rangle (\delta'_1 \wedge pre(e)); \langle <_i \rangle (\delta'_2 \wedge e) \not\vdash \langle \leq_i \rangle \delta'_3$. So now we can use R7 and conclude $\delta'_1;\delta'_2 \not\vdash \delta'_3$. By our induction hypothesis we then get $\delta'_1;\delta'_2 \not\models \delta'_3$, and so Lemma 4.11 delivers $\delta'_1;\delta'_2 \models \neg\delta'_3$. We are shown that for arbitrary δ'_1 and δ'_2. So we have reached our first intermediate goal.

Now we show basically the same as in the previous paragraph, but for a different set of formulas in the first and second argument positions. We want to show that $\delta'_1;\delta'_2 \models \neg\delta'_3$, but this time for any $\delta'_1 \in \Phi_{\delta_1}^{\leq_i} \subseteq S_P^{n-1}$ and $\delta'_2 \in \Phi_{\delta_2}^{\leq_i} \subseteq S_P^{n-1}$. The argument runs just as before, the only difference being that in the case that $\vdash_S \delta'_1 \to pre(e)$, after the applications R3 and R4 we use R8 to conclude $\delta'_1;\delta'_2 \not\vdash \delta'_3$. From there the induction hypothesis and Lemma 4.11 give us our second intermediate goal, namely that $\delta'_1;\delta'_2 \models \neg\delta'_3$ for any $\delta'_1 \in \Phi_{\delta_1}^{\leq_i} \subseteq S_P^{n-1}$ and $\delta'_2 \in \Phi_{\delta_2}^{\leq_i} \subseteq S_P^{n-1}$.

Recall that we proved that $\delta'_1;\delta'_2 \models \neg\delta'_3$ for any $\delta'_1 \in \Phi_{\delta_1}^{\sim_i} \subseteq S_P^{n-1}$ and $\delta'_2 \in \Phi_{\delta_2}^{<_i} \subseteq S_P^{n-1}$. So we get

$$\bigvee \Phi_{\delta_1}^{\sim_i}; \bigvee \Phi_{\delta_2}^{<_i} \models \neg\delta'_3$$

by soundness of R2 and R2′. By the same argument, this time using our second intermediate goal, we conclude that

$$\bigvee \Phi_{\delta_1}^{\leq_i}; \bigvee \Phi_{\delta_2}^{\leq_i} \models \neg\delta'_3$$

So now we are able to use soundness of R6 and the distribution of normal boxes over conjunctions to conclude that

$$[\sim_i]\bigvee \Phi_{\delta_1}^{\sim_i} \wedge [\leq_i]\bigvee \Phi_{\delta_1}^{\leq_i}; [<_i]\bigvee \Phi_{\delta_2}^{<_i} \wedge [\leq_i]\bigvee \Phi_{\delta_2}^{\leq_i} \models [\leq_i]\neg\delta'_3$$

But observe that we have the following

$$\vdash_S \delta_1 \to [\sim_i]\bigvee \Phi_{\delta_1}^{\sim_i} \wedge [\leq_i]\bigvee \Phi_{\delta_1}^{\leq_i}$$

$$\vdash_{\mathcal{E}} \delta_2 \to [<_i](\bigvee \Phi_{\delta_2}^{<_i}) \wedge [\leq_i]\bigvee \Phi_{\delta_2}^{\leq_i}$$

$$\vdash_S \neg\langle \leq_i \rangle \delta'_3 \to \neg\delta_3$$

So we can apply R3 and R4 on the left side and R5 on the right side to conclude that:

$$\delta_1;\delta_2 \models \neg\delta_3$$

from which by Lemma 4.11 we get the required conclusion.

(2b.) $R_i = \leq_i$, and $\delta_1;\delta_2 \nvdash [\leq_i]\bigvee \Phi_i^{\leq_i}$. The proof is essentially the converse of the previous case. Using R3 and R4 we obtain that

$$[\sim_i]\bigvee \Phi_{\delta_1}^{\sim_i} \wedge [\leq_i]\bigvee \Phi_{\delta_1}^{\leq_i}; [<_i]\bigvee \Phi_{\delta_2}^{<_i} \wedge [\leq_i]\bigvee \Phi_{\delta_2}^{\leq_i} \nvdash [\leq_i]\bigvee \Phi_i^{\leq_i}$$

So using R6 we conclude that one of the following must be the case:

$$\bigvee \Phi_{\delta_1}^{\sim_i}; \bigvee \Phi_{\delta_2}^{<_i} \nvdash \bigvee \Phi_i^{\leq_i}$$

$$\bigvee \Phi_{\delta_1}^{\leq_i}; \bigvee \Phi_{\delta_2}^{\leq_i} \nvdash \bigvee \Phi_i^{\leq_i}$$

Let us consider them in turn. By the rule for introduction of $\vee$ on the left side we know that there must be $\delta_1' \in \Phi_{\delta_1}^{\sim_i}$, $\delta_2' \in \Phi_{\delta_2}^{<_i}$ such that:

$$\delta_1'; \delta_2' \nvdash \bigvee \Phi_i^{\leq_i}$$

But we can show something stronger, namely that this holds for δ_1' and δ_2' such that $\vdash \delta_1' \rightarrow pre(e)$ for the unique e such that $\vdash \delta_2' \rightarrow e$. The argument for that is the same as in Aucher's paper. Now we can use our induction hypothesis to get:

$$\delta_1'; \delta_2' \not\models \bigvee \Phi_i^{\leq_i}$$

and using Lemma 4.11 we conclude that

$$\delta_1'; \delta_2' \models \neg \bigvee \Phi_i^{\leq_i}$$

But then (soundness of) R7 gets us

$$\langle \sim_i \rangle(\delta_1' \wedge pre(e)); \langle <_i \rangle(\delta_2 \wedge e)' \models \langle \leq_i \rangle \neg \bigvee \Phi_i^{\leq_i}$$

by pulling out the negation on the right side we get

$$\langle \sim_i \rangle(\delta_1' \wedge pre(e)); \langle <_i \rangle(\delta_2 \wedge e)' \models \neg [\leq_i] \bigvee \Phi_i^{\leq_i}$$

But now observe that $\models \delta_1 \rightarrow \langle \sim_i \rangle(\delta_1' \wedge pre(e))$, $\models \delta_2 \rightarrow \langle <_i \rangle(\delta_2' \wedge e)$ and $\models \neg[\leq_i]\bigvee \Phi_i^{\leq_i} \rightarrow \neg\delta_3$. So using strengthening on the left and weakening on the right we conclude that $\delta_1;\delta_2 \not\models \delta_3$, as desired.
The case where $\bigvee \Phi_{\delta_1}^{\leq_i}; \bigvee \Phi_{\delta_2}^{\leq_i} \nvdash \bigvee \Phi_i^{\leq_i}$ follows the same steps, except we use R8 instead.

■

COROLLARY 4.13 (Soundness and Completeness for $\vdash$)
Let $\phi_1, \phi_3 \in \mathcal{L}_S$ and $\phi_2 \in \mathcal{L}_\mathcal{E}$. Then:

$$\phi_1;\phi_2 \models \phi_3 \text{ iff } \phi_1;\phi_2, \vdash \phi_3$$

PROOF. Same argument as in [1]. ■

5 Generalization and structural properties

This section takes a closer look at valid and invalid structural rules for soft dynamic sequents. This is the second main contribution of this article. To do so we have to generalize soft dynamic sequents so they also cover sequences of updates. Before we present this generalization, however, we recall briefly the structural rules that we shall consider, and explain why the generalization is necessary.

5.1 Tarski consequence relation and soft dynamic sequents

Standard consequence relations satisfy a number of structural conditions. A consequence relation $\Gamma \models \phi$ is a *Tarski consequence relation* whenever it satisfies the following conditions, for any finite set of formula Γ, Σ and formulas ϕ, ψ and χ:[7]

- if $\Gamma, \phi, \phi \models \psi$, then $\Gamma, \phi \models \psi$ [contraction]
- if $\Gamma, \phi, \psi, \Sigma \models \chi$, then $\Gamma, \psi, \phi, \Sigma \models \chi$ [permutation, a.k.a. exchange]
- $\phi \models \phi$ [reflexivity, a.k.a. inclusion]
- if $\Gamma \models \phi$, then $\Gamma, \psi \models \phi$ [monotonicity, a.k.a. weakening]
- if $\Gamma \models \phi$ and $\Sigma, \phi \models \psi$, then $\Gamma, \Sigma \models \psi$ [transitivity, a.k.a. cut]

The set of valid soft dynamic sequents does not constitute a Tarski consequence relation. Soft dynamic sequents are triples of single formulae (a state formula, an event formula and a state formula). The result is that there is no way to formulate contraction, and no *straightforward* way to formulate exchange. Some form of exchange might be valid modulo with a pair of translations $\tau_1 : \mathcal{L}_{\mathcal{E}} \to \mathcal{L}_S$, $\tau_2 : \mathcal{L}_S \to \mathcal{L}_{\mathcal{E}}$:

- if $\phi; \psi \models \chi$, then $\tau_1(\psi); \tau_2(\phi) \models \chi$

But lexicographic update ignores the valuation of atomic propositions in event models, so we conjecture that exchange will turn out to be invalid for any meaningful way to spell out τ_1 and τ_2. Reflexivity can be formulated for state formulas.

- $\phi_1; \phi_2 \models \phi_1$

But this is invalidated by the same kind of 'Moore sentences' that cause failure of reflexivity for dynamic consequences for public announcements [21].

As mentioned above, monotonicity or weakening is also not expressible in soft dynamic sequents, but rules like R3 capture it through reference to the logical strength in the logic of plausibility and event models. Conceptually this can be seen as a surprise. Soft attitudes like conditional beliefs are non-monotonic. But the semantic definition of valid soft dynamic sequents is a universal quantification: for all $\mathcal{M}, w$ and $\mathcal{E}, e$, if $\mathcal{M}, w \models \phi_1$ and $\mathcal{E}, e \models \phi_3$, then $(\mathcal{M} \otimes \mathcal{E}, (w, e)) \models \phi_3$. Such universal quantifications are monotonic.

Finally, there is a version of cut that can be formulated in soft dynamic sequents:

- if $\phi_1; \phi_2 \models \phi_3$ and $\phi_3; \phi_2 \models \psi$, then $\phi_1; \phi_2 \models \psi$ [contractive cut]

The rule is "contractive" in that it not only cuts the middle term ϕ_3, it also suppresses one occurrence of the ϕ_2-type of update.

[7] The permutation rule below is strong enough to entail associativity. Although we do not explore it here, a more fine-grained approach would involve a restricted permutation rule together with an explicit rule for associativity. We owe this point to an anonymous referee.

FACT 5.1
Contractive cut is *not* valid for soft dynamic sequents

PROOF. (Sketch) Counterexamples are easy to devise using the following observation. From the first premise we know that after an update by a ϕ_2-type of epistemic action in a ϕ_1-initial situation we always end up in a situation where ϕ_3 holds. But to use the second premise and reach a ψ-type of situation, as required to establish the consequent of the rule, we need an *additional* update: a ϕ_2-type of epistemic action. There is no certainty that the first update will suffice. ■

The failure of contractive cut shows why a generalization to sequences of updates is indeed of interest. As we shall see below, multiplicative cut, viz., cut that brings with it the second occurrence of ϕ_2, is valid for generalized soft dynamic sequents. Generalizing to sequences of updates really shows the sub-structural character of soft dynamic sequents.

5.2 Generalized soft dynamic sequents

We want to generalize soft dynamic sequents in such a way that the structural rules just mentioned can be expressed more naturally. To do this we allow for sequences of update formulas in the second argument. Conceptually and formally, this is the important difference from the sequents studied so far. We also allow finite sets of formulas to occur in the first argument. This, of course, is equivalent to what we had before, viz., a single formula in the first argument. We do this to make generalized soft dynamic sequents notationally closer to consequence relations, as studied in substructural logic.

DEFINITION 5.1 (Dynamic Consequence Relation – Generalized)
Let Γ be a finite set of formulas of $\mathcal{L}_S$ and Σ be a finite sequence $(\psi_1,...,\psi_n)$ of formulas of $\mathcal{L}_\mathcal{E}$. A *generalized soft dynamic sequent* (GSDS) is an object of the form:

$$\Gamma;\Sigma \models \phi$$

A GSDS is *valid* iff for all M,w and sequence $((\mathcal{E}_1,e_1),...(\mathcal{E}_n,e_n))$ of length n, if $M,w \models \bigwedge\Gamma$ and for all $1 \leq i \leq n$ we have that $E_i,e_i \models \psi_i$ then $((\mathcal{M}\otimes\mathcal{E}_1)\otimes...)\otimes\mathcal{E}_n),((w,e_1)...,e_n) \models \phi_3$.

Generalized soft dynamic sequents are close to the consequence relations investigated by [21] and [8] in that they now cover sequences of updates, and are contextualized. But in contrast to [8], we keep to finite descriptions of the initial situation, i.e. finite set Γ in the first argument.

Each of the structural rules mentioned above can now be formulated component-wise, i.e. either for state or event formulas. For instance, one can formulate two versions of exchange, here in simple, two-formula form:

- if $\{\phi,\psi\};\Sigma \models \chi$, then $\{\psi,\phi\};\Sigma \models \chi$ [exchange, state formulas]
- if $\Gamma;\langle\phi,\psi\rangle \models \chi$, then $\Gamma;\langle\psi,\phi\rangle \models \chi$ [exchange, event formulas]

But, as before, there is no counterpart of these rules that crosses the semicolon. The distinction between state and event formulas remains intact. Structural rules operate within each context, i.e. to describe the situations the agents are in and to describe the types of epistemic actions that are occurring, but not across contexts.

5.3 Structural properties of soft dynamic sequents

Although the dynamic relations are not Tarski relations, they do have a number of weaker structural properties. Here we offer a brief structural analysis of the dynamic consequence relation for lexicographic update.

In a sequent $\Gamma, \Sigma \models \phi$ the first position will now immediately have contraction and permutation as admissible rules. This follows trivially from Γ being a set. In addition, the consequence relation becomes monotonic for the first position.

Fact 5.2
The following is valid for GSDS:

$$\Gamma; \Sigma \models \phi, \text{ then } \Gamma, \psi; \Sigma \models \phi$$

Proof. Analogous to soundness of R3 in for soft dynamic sequents. ■

Generalized soft dynamic sequents are much less structural in the second position, i.e. for sequences of update. Monotonicity, exchange and contraction all fail there. Failure of exchange within sequences Σ arises for reasons similar to those given by [21] in the case of public announcements. For a counter-example, observe that $\{\neg pre(e)\}; \langle e, [<_i]\mathrm{P}\rangle \models \perp$ is a valid GSDS, but that $\{\neg pre(e)\}; \langle [<_i]\mathrm{P}, e\rangle \models \perp$ is not. One can easily devise a case where updating first by a $[<_i]\mathrm{P}$ makes e-types of epistemic actions—i.e. any action with precondition $pre(e)$—executable. This failure of exchange reflects the fact that soft information updates are order-dependent. This dependency was already well known for hard information updates, see again [21]. But for general information dynamics, this is aggravated by the fact that epistemic actions can have arbitrary conditions of executability, as the counter-example illustrates.

Failure of contraction reflects the fact that iterated update with the same type of epistemic action is *not* reducible to a single update of that type. This is also a well-known fact in the theory of multi-agent-iterated belief revision, see e.g. [4, 18].

There are two forms of monotonicity in the second position, only one of which fails. First, there is the straightforward generalization of R4:

Fact 5.3
The following monotonicity rule is valid for soft information sequents:

$$\frac{\Gamma; \Sigma \vdash \phi}{\Gamma; \Sigma[\phi_i/\psi_i] \vdash \phi}\ R4{-}Gen$$

with $\vdash \psi_i \rightarrow \phi_i$ and $\Sigma[\phi_i/\psi_i]$ being the sequence resulting from substituting ψ_i for the ith component ϕ_i in Σ.

Proof (Sketch). Suppose $\Gamma; \Sigma \vdash \phi$. Take a pointed plausibility model that satisfies all the formulas in Γ. Start the sequence of update $\Sigma[\phi_i/\psi_i]$. This sequence is the same as Σ up to the ith step. There one updates the current plausibility model with a ψ_i type of (pointed) event models instead of a ϕ_i. But because $\vdash \psi_i \rightarrow \phi_i$ we know that this pointed event model also satisfies ϕ_i. But then, since Σ and $\Sigma[\phi_i/\psi_i]$ agree on subsequent updates, we know from $\Gamma; \Sigma \vdash \phi$ that ϕ will be true in the pointed plausibility model resulting from updating with $\Sigma[\phi_i/\psi_i]$.

■

With sequences of formulas in the second position we can formulate a more classical form of monotonicity, namely one that adds formulas to the existing sequence. By doing so we do not move from a logically weaker to logically stronger antecedent. This type of operation adds one update in the sequence.

FACT 5.4
The following monotonicity rules are invalid for soft information sequents:

$$\frac{\Gamma;\Sigma\vdash\phi}{\Gamma;\psi,\Sigma\vdash\phi}\ LM \quad \frac{\Gamma;\Sigma\vdash\phi}{\Gamma;\Sigma,\psi\vdash\phi}\ RM$$

PROOF (SKETCH). Right monotonicity is not valid because updating with an additional ψ-type of epistemic action can break the conclusion of the sequence of update Σ. Again, examples are easy to come by. Failure of left monotonicity for contextualized dynamic consequence has already been observed by [8]. Their counter-example can be easily adapted to lexicographic upgrade. ■

In Section 5.1, we mentioned that contractive cut is not valid for soft dynamic sequents. We observed that that rule failed because it contracted two ϕ_2-updates into one. The present generalization allows us to formulate a more careful, multiplicative version of cut, which turns out to be valid.

FACT 5.5
The following multiplicative cut rule is valid for soft information sequents:

$$\frac{\Gamma;\Sigma\vdash\phi \quad \{\phi\};\Sigma'\vdash\psi}{\Gamma;\Sigma,\Sigma'\vdash\psi}\ MCut$$

PROOF (SKETCH). Take a pointed plausibility model that satisfies all the formulas in Γ. By the left-most premise, applying the sequence of updates Σ will result in a pointed plausibility model that satisfies ϕ. But by the right-most premise, continuing the update sequences with Σ' will result in a model where ψ holds. ■

The reader can check that the natural generalization of most of the rules in Table 1 are valid, even for the box introduction rules R6 and R9. Rules for diamond introduction are more intricate to devise with sequences of update. One needs to refer to the preconditions of the updates after the first one in a sequence Σ. All the rules that we have considered at this point fail to do this in a valid way. The validity of multiplicative cut also shows that this generalization of soft dynamic sequents validates some rules that allow sequences of updates to be extended. Multiplicative cut concatenates sequences. It is still an open question whether there are other rules that allow one to operate on sequences in a similar fashion. For these reasons we do not know whether the set of valid generalized soft dynamic sequents is completely axiomatizable.

6 Conclusion

This article has extended the study of dynamic consequence relations to a general class of soft information changes. We have provided a sound and complete calculus for one-step soft dynamic consequence relations, and given a number of valid and invalid structural rules for a generalization to sequences of updates. The complete axiomatization of the latter has been left open: this is the obvious next step for future work. It would also be interesting to look at dynamic consequences for specific update rules for soft information that have been studied in the literature, such as radical and conservative upgrade, in particular whether they validate additional structural rules.

Acknowledgments

The authors would like to thank the participants of the Prague Workshop on Sub-Structural Epistemic Logic (June 2011), and of the Munich Meeting on Sub-structural Epistemic Logic (January 2013) for very useful comments and suggestions on earlier versions of the article. They are also very grateful to the two anonymous referees of the *JLC* for thorough reading and extensive comments.

Part of this work has been done while O.R. was affiliated at the LMU Munich. Financial support of this institution, and of the Alexander von Humboldt Foundation, is gratefully acknowledged.

References

[1] G. Aucher. DEL-sequents for progression. *Journal of Applied Non-Classical Logics*, **21**, 289–321, 2011.

[2] G. Aucher. Update Logic. Research Report RR-8341, INRIA, August 2013. http://hal.inria.fr/hal-00849856.

[3] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In *vol. 3 of Texts in Logic and Games*, G. Boranno, W. van der Hoek and M. Wolldridgge (eds), pp. 13–60. Amsterdam University Press, 2008.

[4] A. Baltag and S. Smets. Group belief dynamics under iterated revision: fixed points and cycles of joint upgrades. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, A. Heifetz, ed, pp. 41–50. ACM, 2009.

[5] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *TARK 98*, 1998.

[6] A. Baltag, S. Smets, and J. Zvesper. Keep 'hoping' for rationality: a solution to the backward induction paradox. *Synthese*, **169**, 301–333, 2009.

[7] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2002.

[8] A. Cordón-Franco, H. van Ditmarsch, and A. Nepomuceno-Fernández. Dynamic consequences and announcement. *Review of Symbolic Logic*, **6**, 659–679, 2013.

[9] C. Dégremont. *The Temporal Mind: Observations On the Logic of Belief Change in Interactive Systems*. PhD Thesis, ILLC, University of Amsterdam, 2010.

[10] F. Dietrich, C. List, and R. Bradley. A joint characterization of belief revision rules. MPRA Paper 41240, University Library of Munich, Germany, September 2012. http://ideas.repec.org/p/pra/mprapa/41240.html.

[11] K. Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, **16**, 229–237, 1975.

[12] P. Gärdenfors. Belief revision and nonmonotonic logic: two sides of the same coin? *Logics in AI*, vol. 478, pp. 52–54. LNCS, 1991.

[13] G. Greco, A. Kurz, and A. Palmigiano. Dynamic epistemic logic displayed. In *Logic, Rationality and Interaction, Proceedings of the 4th International Workshop*. D. Grossi, O. Roy, and H. Huang, eds, LNCS, Folli Series, 2013.

[14] P. Maffezioli and S. Negri. A Gentzen-style analysis of public announcement logic. In *Proceedings of the International Workshop on Logic and Philosophy of Knowledge, Communication and Action*, pp. 293–313, 2010.

[15] D. Makinson. *Bridges from Classical to Nonmonotonic Logic*. King's College, 2005.

[16] M. Martinez and S. Sequoiah-Grayson. Logic and information. In *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta, ed, Spring 2014 edn, 2014.

[17] L. S. Moss. Finite models constructed from canonical formulas. *Journal of Philosophical Logic*, **36**, 605–640, 2007.

[18] E. Pacuit and O. Roy. A dynamic analysis of interactive rationality. In *Logic, Rationality, and Interaction*, H. van Ditmarsch, J. Lang, and S. Ju, eds, pp. 244–257. Springer, 2011.

[19] G. Restall. *An Introduction to Sub-Structural Logics*. Routledge, London, 2000.

[20] J. van Benthem. *Exploring Logical Dynamics*. CSLI publications Stanford, 1996.

[21] J. van Benthem. Logical dynamics meets logical pluralism? *The Australasian Journal of Logic*, **6**, 182–209, 2008.

[22] J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.

[23] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, **204**, 1620–1662, 2006.

[24] J. van Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, **38**, 491–526, 2009.

[25] J. van Benthem, P. Girard, and O. Roy. Everything else being equal: a modal logic for *ceteris paribus* preferences. *Journal of Philosophical Logic*, **38**, 83–125, 2009.

[26] H. van Ditmarsch, W. van de Hoek, and B. Kooi. *Dynamic Epistemic Logic*, vol. 337 of *Synthese Library Series*, Springer, 2007.

[27] F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, **25**, 221–261, 1996.

[28] Y. Venema. *Modal Logic, Algebra and Coalgebra*. ILLC, Amsterdam, 2012. Downloaded from the author's website on 01-03-12.

Received 20 December 2013