

Dynamic consequence and public announcement

Andrés Córdón Franco*, Hans van Ditmarsch†, Angel Nepomuceno‡

August 2, 2013

Abstract

In [25] van Benthem proposes a dynamic consequence relation defined as

$$\psi_1, \dots, \psi_n \models^d \varphi \text{ iff } \models^{pa} [\psi_1] \dots [\psi_n] \varphi,$$

where the latter denotes consequence in public announcement logic, a dynamic epistemic logic. In this paper we investigate the structural properties of a conditional dynamic consequence relation \models_{Γ}^d extending van Benthem's proposal. It takes into account a set of background conditions Γ , inspired by [18] wherein Makinson calls this reasoning 'modulo' a set Γ . In the presence of common knowledge, conditional dynamic consequence is definable from (unconditional) dynamic consequence. An open question is whether dynamic consequence is compact. We further investigate a dynamic consequence relation for soft instead of hard announcements. Surprisingly, it shares many properties with (hard) dynamic consequence. Dynamic consequence relations provide a novel perspective on reasoning about protocols in multi-agent systems.

1 Introduction and preliminaries

In dynamic semantics or *update semantics*, dynamic consequence relations are an obligatory phenomenon and are well-studied [29, 23, 15, 17]. Sometime in the late 1980s, *dynamic epistemic logic* [28] parted ways with dynamic semantics and put all the dynamics in modal dynamic operators instead. These modal operators are interpreted employing a standard Kripke semantics. The consequence relation associated with this semantics has nothing dynamic. In [25], van Benthem proposes a way that goes back from dynamic epistemic logic to dynamic consequence. He defines a dynamic consequence relation as: $\psi_1, \dots, \psi_n \models^d \varphi$

*Department of Computer Science and Artificial Intelligence, Universidad de Sevilla, Av. Reina Mercedes s/n, 41012 Seville, Spain. Email: acordon@us.es

†LORIA, CNRS – Université de Lorraine, BP 239, 54506 Vandoeuvre-lès-Nancy, France. Email: hans.van-ditmarsch@loria.fr

‡Department of Philosophy and Logic and Philosophy of Science, Universidad de Sevilla, Av. Camilo José Cela s/n, 41018 Seville, Spain. Email: nepomuce@us.es

iff $\models^{pa} [\psi_1] \dots [\psi_n] \varphi$, where the latter denotes consequence in public announcement logic, a dynamic epistemic logic. In public announcement logic, $[\psi]\varphi$ stands for ‘after (truthful public) announcement of ψ , φ (is true)’.

Example 1 *After public announcement of factual information p (‘It rains in Seville’), the agent knows that p (‘Fernando knows that it rains in Seville’). This is formalized as $[p]Kp$. Seen as consequence, we then have that $p \models^d Kp$. This new form of dynamic consequence is not trivial. For (another) example, observe that $p \wedge \neg Kp \not\models^d p \wedge \neg Kp$: if I inform you that (p is true and you do not know that p), then you know p as a result. (‘Fernando calls Graham, who lives in Australia, and tells him: “You don’t know that it is raining in Seville!”.’ With the usual conversational implicature, this is an announcement of $p \wedge \neg Kp$. Then, of course, Graham knows p , so $p \wedge \neg Kp$ is now false.)*

We propose a dynamic consequence relation \models_{Γ}^d , called *conditional dynamic consequence*, that extends van Benthem’s proposal. It is an extension of \models^d because it also takes a third parameter into account in the consequence relation, namely a set of *background conditions* Γ . This is a set of formulas commonly known by the agents. The reason to propose such an extension is that the study of the consequences of sequences of announcements is often in the context of a given system satisfying a number of initial properties. Now there is nothing per se that prevents that such properties are made known to all agents by yet another sequence of dynamic phenomena—initializing the system, one might say, from some sort of perspective of common ignorance; but in fact such initial configurations are often a *given*. Also, in such initial configurations agents have different perspectives. Such different perspectives cannot have been the consequence of making public announcements in a state of commonly known ignorance, but need more complex dynamics, such as private announcements. A more elegant starting point therefore seems the assumption of that third parameter. This assumption also matches the general observation on multi-S5 systems that the structure of the model is commonly known to all agents, and that for many such systems (for example, those with finite models) this structure can be captured by a set of commonly known formulas. This is the background theory, the set of background conditions.

We then investigate the structural properties of conditional dynamic consequence. Like van Benthem’s dynamic consequence, the structural rules permutation, contraction, reflexivity, cut and monotonicity fail. Van Benthem’s dynamic consequence satisfies cautious monotonicity, left cut, and left monotonicity (these structural rules will later be defined). Conditional dynamic consequence satisfies cautious monotonicity and left cut, but does not satisfy left monotonicity.

We further investigate how the presence of common knowledge in the logical language influences the results, and give tentative results on what the consequences are of allowing infinite sets of premises, in relation to a suitable notion of compactness. We also present a version of the dynamic consequence relation for so-called *soft* public announcements, for which the results are (surprisingly) similar to those for (standard, or ‘hard’) public announcements.

1.1 Public announcement logic

Dynamic epistemic logics [28] combine modal operators for knowledge or belief, typically for a set of agents, with dynamic modal operators for change of knowledge or belief. Public announcement logic [19, 8] is one such dynamic epistemic logic, namely where the dynamics consists of public truthful announcements (unlike, e.g., private announcements to subgroups). We present the version of the logic with common knowledge operators.

The *language of public announcement logic* [19] is defined over a finite set of agents A and a countable set of *primitive propositions* (or *propositional variables*) P . Let $a \in A$ and $p \in P$:

$$\mathcal{L}_C^{pa} \quad \ni \quad \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid [\varphi]\varphi \mid C\varphi$$

Other propositional connectives and the constants \top ('true') and \perp ('false') are defined by the usual abbreviations. For arbitrary formulas we write φ, ψ, \dots . Formula $K_a\varphi$ stands for 'agent a knows φ ', $C\varphi$ for 'the agents commonly know φ ' and $[\varphi]\psi$ for 'after public announcement of φ, ψ '. Sets of formulas are represented by Γ, Γ', \dots or Δ, \dots and sequences of formulas by Σ, \dots (for *Sequence*).

We use various sublanguages in our presentation. Without the inductive construct for common knowledge, we get the language \mathcal{L}^{pa} , without public announcement as well we get the language of (multi-agent) epistemic logic \mathcal{L}^{el} , and \mathcal{L}_C^{el} is the language with common knowledge but without public announcement. The language \mathcal{L}^{pl} of propositional logic, the booleans, does not have any modalities. The basis for our presentation of dynamic consequence will be \mathcal{L}^{el} . The reason for this restriction is that we try to obtain our results for a minimal setting, and generalize the results from there. With common knowledge we get different results.

A *Kripke structure* or *epistemic model* over A and P is a tuple $M = (S, \{\sim_a\}_{a \in A}, V)$ where S is a (non-empty) set or domain of states, for each agent $a \in A$, $\sim_a \subseteq S \times S$ is an epistemic accessibility or indistinguishability relation that is assumed to be an equivalence relation, and valuation $V : P \rightarrow \mathcal{P}(S)$ assigns primitive propositions to the set of states in which they are true. For 'state s is in the domain of model M ' we write $s \in M$. A *pointed Kripke structure* is a pair (M, s) where $s \in M$ (we often omit the two parentheses in constructs wherein (M, s) occurs). To interpret common knowledge we further need the reflexive transitive closure of the union of the accessibility relations of the agents, namely $\sim_A ::= (\bigcup_{a \in A} \sim_a)^*$. This is also an equivalence relation. We write $K_a\Gamma$ for $\{K_a\varphi \mid \varphi \in \Gamma\}$ and similarly for $C\Gamma$.

The interpretation of the crucial constructs in a pointed Kripke structure is defined as follows—in this definition, the symbol \models^{pa} is more properly a forcing and not a logical consequence symbol, but the convention in epistemic logic is to allow this overload of notation, as the two coincide in the notion of validity.

$$M, s \models^{pa} p \text{ iff } p \in V(p)$$

$$M, s \models^{pa} \varphi \wedge \psi \text{ iff } M, s \models^{pa} \varphi \text{ and } M, s \models^{pa} \psi$$

$$M, s \models^{pa} \neg\varphi \text{ iff } M, s \not\models^{pa} \varphi$$

- $M, s \models^{pa} K_a \varphi$ iff for every t such that $s \sim_a t$, $M, t \models^{pa} \varphi$
- $M, s \models^{pa} C \varphi$ iff for every t such that $s \sim_A t$, $M, t \models^{pa} \varphi$
- $M, s \models^{pa} [\varphi] \psi$ iff $M, s \models^{pa} \varphi$ implies that $M|_{\varphi}, s \models^{pa} \psi$

where $M|_{\varphi} = (S', \{\sim'_a\}_{a \in A}, V')$ such that $S' = \{s' \in S : M, s' \models^{pa} \varphi\}$; $\sim'_a = \sim_a \cap (S' \times S')$; $V'(p) = V(p) \cap S'$. We write $M, s \models^{pa} \Gamma$ for (for all $\varphi \in \Gamma$, $M, s \models^{pa} \varphi$).

We further define that ‘ φ is valid on M ’, notation $M \models^{pa} \varphi$, as ‘for all $s \in M$, $M, s \models^{pa} \varphi$ ’, and ‘ φ is valid’ (or *pa-valid*), notation $\models^{pa} \varphi$, as ‘for all M (in the class of models for agents in A and primitive propositions in P), φ is valid on M ’. Logical consequence \models^{pa} is defined as: $\Gamma \models^{pa} \varphi$ iff for all models M and states $s \in M$, if $M, s \models^{pa} \psi$ for all $\psi \in \Gamma$, then $M, s \models^{pa} \varphi$. We write $\Delta \models^{pa} \Gamma$ for: $\Delta \models^{pa} \psi$ for all $\psi \in \Gamma$.

1.2 Structural properties of consequence relations

Assume a logical language \mathcal{L} . A *consequence relation* is a binary relation \models between a sequence Σ of formulas in \mathcal{L} , the *premises*, and a formula in \mathcal{L} , the *conclusion*. An example is the just introduced \models^{pa} ; where the premises constitute a set, not a sequence. Consequence relations may satisfy certain *structural rules*, the standard rules are:

- $\varphi \models \varphi$ *reflexivity*
- $\Sigma, \varphi, \varphi \models \psi$ implies $\Sigma, \varphi \models \psi$ *contraction*
- $\Sigma, \varphi, \varphi', \Sigma' \models \psi$ implies $\Sigma, \varphi', \varphi, \Sigma' \models \psi$ *permutation*
- $\Sigma \models \psi$ and $\Sigma', \psi, \Sigma'' \models \varphi$ imply $\Sigma', \Sigma, \Sigma'' \models \varphi$ *cut*
- $\Sigma \models \psi$ implies $\Sigma, \varphi \models \psi$ *monotonicity*

For relational studies of logical consequence, see [22, 11]. For an overview of various logics satisfying fewer or different structural rules, see e.g. [20]. The modified structural rules considered by van Benthem in [25] for dynamic consequence are as follows.

- $\Sigma \models \psi$ and $\Sigma, \psi, \Sigma' \models \varphi$ imply $\Sigma, \Sigma' \models \varphi$ *left cut*
- $\Sigma \models \psi$ implies $\varphi, \Sigma \models \psi$ *left monotonicity*
- $\Sigma \models \psi$ and $\Sigma, \Sigma' \models \varphi$ imply $\Sigma, \psi, \Sigma' \models \varphi$ *cautious monotonicity*

We are now sufficiently prepared to present van Benthem’s dynamic consequence relation and our conditional dynamic consequence relation.

2 Conditional dynamic consequence

Using the standard Kripke semantics of public announcement logic one can define a dynamic consequence relation where the premises stand for consecutive announcements, and the conclusion stands for a postcondition of these announcements. Given a set of agents A and a set of atoms P , let $\Sigma = \psi_1, \dots, \psi_n$ be a (finite!) sequence of \mathcal{L}^{el} formulas, and $\varphi \in \mathcal{L}^{el}$. We write $[\Sigma]\varphi$ for $[\psi_1] \dots [\psi_n]\varphi$ (we note that $[\Sigma]\varphi \in \mathcal{L}^{pa}$), and we write $M|\Sigma$ for $(\dots (M|\psi_1)| \dots)|\psi_n$.

Definition 2 (Dynamic consequence [25])

$\Sigma \models^d \varphi$ iff $\models^{pa} [\Sigma]\varphi$.

If $\Sigma \models^d \varphi$ we say that φ is a *dynamic consequence* of Σ , or Σ dynamically entails φ . In other words, Σ dynamically entails φ iff for every model M and state $s \in M$, it holds that $M, s \models^{pa} [\psi_1] \dots [\psi_n]\varphi$; where the latter means that on condition that ψ_1 is true in (M, s) , and ψ_2 is true in $(M|\psi_1, s)$, and so on, then we have that $M|\Sigma, s \models^{pa} \varphi$. We emphasize that, as this should be the case for every state s in M satisfying the announcement sequence, φ is therefore not merely true in some state of $M|\Sigma$ but a model validity on $M|\Sigma$.¹

A main result in [25] is that \models^d satisfies the structural rules of left monotonicity, left cut, and cautious monotonicity, that were introduced in the previous section²; and that these three rules even *represent* that consequence relation, where ‘represent’ is interpreted in a technical sense lifting the discussion to abstract transition models, see [25, pp.193-194] for details.

We now extend this consequence relation. Let Γ be a set of \mathcal{L}^{el} formulas, called the set of *background conditions* of the consequence relation. (Γ may be infinite.)³

Definition 3 (Conditional dynamic consequence)

$\Sigma \models_{\Gamma}^d \varphi$ iff for all M , $M \models^{pa} \Gamma$ implies $M \models^{pa} [\Sigma]\varphi$.

We say that Σ *dynamically entails φ conditional to background Γ* . Obviously, \models^d is the special case of \models_{Γ}^d for $\Gamma = \{\top\}$. If Γ is a singleton set $\{\varphi\}$ we write \models_{φ}^d instead of $\models_{\{\varphi\}}^d$.⁴

¹Van Benthem does not really define dynamic consequence as $\Sigma \models^d \varphi$ iff $\models^{pa} [\Sigma]\varphi$, but as $\Sigma \models^d \varphi$ iff $\models^{pa} [\Sigma]C\varphi$. In the language \mathcal{L}_C^{pa} of public announcement logic with common knowledge this amounts to the same. In that logic, $\models^{pa} [\psi]\varphi$ is equivalent to $\models^{pa} [\psi]C\varphi$, for all φ and ψ . (If $[\psi]\varphi$ is valid, then in any given model, φ is valid on its ψ -restriction. But in that case $C\varphi$ is also valid on the ψ -restriction. Therefore $[\psi]C\varphi$ is valid on that model. As the model was arbitrary, $[\psi]C\varphi$ is valid. The other direction is even simpler.) And therefore we also have this for sequences of announcements.

²This does not depend on the presence of common knowledge in van Benthem’s setting. See Section 3.

³Our inspiration was the so-called *pivotal-assumption consequence* in [18]. Makinson proposes logical consequence *modulo* an *assumption set* K . This assumption set K in [18] plays the role of our Γ . Otherwise, his setting is different. E.g., Makinson’s has classical structural properties, and the pivotal assumptions / background conditions Γ can be added to the premises, unlike in our proposal.

⁴In a prior version [10] of this work we distinguished conditional dynamic consequence (there called global dynamic consequence) from *local dynamic consequence*, defined as

$$\Sigma \models_{\Gamma}^{ld} \varphi \text{ iff } \Gamma \models^{pa} [\Sigma]\varphi.$$

Example 4 Consider a pack of cards dealt over different players, such that every player only knows his own card. An epistemic model encoding this uncertainty of the players about the card deal can be described by a set of formulas [27]. This set consists of the available **facts**, the available (positive) **knowledge**, and the available **ignorance**. The facts are, that one of all possible deals of cards is actually the case. Each player knows his own hand of cards. And each player considers all deals of cards possible wherein she holds her own cards (this entails that each player is ignorant of the ownership of all other cards than her own). For three players and three cards, with therefore nine propositions c_a for ‘player a holds card c ’, a valuation (‘the facts’) describes a card deal (like $0_a \wedge 1_b \wedge 2_c \wedge \neg 1_a \wedge \neg 2_a \wedge \dots$ —four more negated variables), knowledge is encoded by propositions like $0_a \rightarrow K_a 0_a$ (if player a holds card 0, she knows that), and ignorance by propositions like $\neg K_a 2_b$ (player a does not know that player b holds card 2).

Let this theory describing a card deal model be called T . Typical sequences of announcements that one wishes to check for their consequences in this setting, are protocols wherein the two players a, b can achieve knowledge of the card deal by means of public announcements without the third player c getting to know the card deal. These protocols Σ should satisfy information requirements φ (containing components such as $2_b \rightarrow K_a 2_b$, ‘if b holds card 2, a knows it’) and safety requirements ψ (containing components such as $2_b \rightarrow \neg K_c 2_b$, ‘... but c does not’). The question is then whether

$$\Sigma \models_T^d \varphi \wedge \psi.$$

Example 5 We have that $p \models_{p \wedge \neg K p}^d \neg p$, whereas $p \wedge \neg K p, p \not\models^d \neg p$. The first holds because $p \wedge \neg K p$ cannot be a model validity (in $S5$).

Example 5 shows that background conditions cannot, instead, be simply announced first. Section 2.4 will discuss when they can and when they cannot.

2.1 The failure of classical structural rules

The so-called classical structural rules reflexivity, contraction, permutation, monotonicity, and cut all fail for conditional dynamic consequence. In [25] it is shown that they fail for

Although on first sight the more obvious definition, it does not to capture the intuition that the formulas in Γ are background conditions. For example, we then have that

$$\neg K \neg p \models_p^{ld} p.$$

It is true that ‘on condition that p is true, after announcement of $\neg K \neg p$, p is then always true’, because propositional variables do not change their value in public announcement logic. But we cannot then say that ‘ p is background knowledge for the agent’. If that were so, we would expect p to be known by the agent, no matter what the subsequent announcement is. But of course,

$$\neg K \neg p \not\models_p^{ld} K p.$$

(The example does not need a premise: we also have that $\models_p^{ld} p$ but $\not\models_p^{ld} K p$.) That we cannot strengthen the consequent seems a strong argument against this local dynamic consequence. Local dynamic consequence satisfies even fewer structural rules than conditional dynamic consequence.

the dynamic consequence relation \models^d without background, which is the borderline case for $\Gamma = \{\top\}$ of \models_Γ^d . It is instructive to show counterexamples.

- *Reflexivity fails.* As known, the formula

$$[p \wedge \neg K_a p] \neg (p \wedge \neg K_a p)$$

is valid, so

$$p \wedge \neg K_a p \not\models^d p \wedge \neg K_a p$$

- *Contraction fails.* We again make use of the Moore-sentence $p \wedge \neg K_a p$ to show that.

– $p \wedge \neg K_a p, p \wedge \neg K_a p \models^d \perp$, because $p \wedge \neg K_a p$ cannot be truthfully announced more than once.

– $p \wedge \neg K_a p \not\models^d \perp$.

- *Permutation fails.*

– $p, \neg K_a p \models^d \perp$, because after announcing p , this is known, so $\neg K_a p$ cannot be truthfully announced after that.

– $\neg K_a p, p \not\models^d \perp$, because after announcing your ignorance of a fact, this fact remains unknown; so it can then be very well announced truthfully.

- *Monotonicity fails.* A counterexample is as follows.

$$\neg K_a p \models^d \neg K_a p,$$

however,

$$\neg K_a p, p \not\models^d \neg K_a p,$$

since

$$\not\models^{pa} [\neg K_a p][p] \neg K_a p.$$

- *Cut fails.* Take (see the definition of cut) $\Sigma = \neg p$, $\psi = K_a \neg p$, $\Sigma' = \neg K_a \neg p$, $\Sigma'' = \emptyset$, and $\varphi = \perp$. Then it holds that

– $\neg p \models^d K_a \neg p$, and

– $\neg K_a \neg p, K_a \neg p \models^d \perp$; but

– $\neg K_a \neg p, \neg p \not\models^d \perp$.

2.2 Structural rules for conditional dynamic consequence

In view of the failure of classical structural rules, it is natural to ask ourselves if conditional dynamic consequence satisfies any variant of these rules. The prime candidates are the modified structural rules ‘left monotonicity’, ‘left cut’, and ‘cautious monotonicity’, that are satisfied by unconditional dynamic consequence. The first one is not satisfied but the other two are. We prove the results for conditional but also for unconditional dynamic consequence, as [25] does not give proofs. This we will address first. But we can also consider rules for changing the background conditions Γ of our consequence relation \models_{Γ}^d , or consider special background conditions. That will be addressed subsequently.

We need a simple lemma, used in the theorem that follows it.

Lemma 6 *If $M \models^{pa} [\Sigma]\varphi$, then for all χ , $M \models^{pa} [\Sigma][\varphi]\chi \leftrightarrow [\Sigma]\chi$.*

Proof The assumption $M \models^{pa} [\Sigma]\varphi$ entails that φ is a model validity on $M|\Sigma$. Therefore, announcing it in model $M|\Sigma$ has no informative effect: it is the trivial model restriction. \square

Theorem 7 *Conditional dynamic consequence satisfies cautious monotonicity and left cut but not left monotonicity.*

Proof

- *Cautious monotonicity*

In order to show that $\Sigma, \varphi, \Sigma' \models_{\Gamma}^d \psi$ we assume that $\Sigma \models_{\Gamma}^d \varphi$ and that $\Sigma, \Sigma' \models_{\Gamma}^d \psi$. Let model M be such that $M \models \Gamma$. From $M \models \Gamma$ and assumption $\Sigma \models_{\Gamma}^d \varphi$ follows $M \models^{pa} [\Sigma]\varphi$. From $M \models \Gamma$ and assumption $\Sigma, \Sigma' \models_{\Gamma}^d \psi$ follows $M \models^{pa} [\Sigma][\Sigma']\psi$. From $M \models^{pa} [\Sigma]\varphi$, $M \models^{pa} [\Sigma][\Sigma']\psi$ and Lemma 6, by taking $\chi = [\Sigma']\psi$, follows $M \models^{pa} [\Sigma][\varphi][\Sigma']\psi$.

- *Left cut*

It suffices to observe that on assumption of $M \models^{pa} [\Sigma]\varphi$, from $M \models^{pa} [\Sigma][\varphi][\Sigma']\psi$ also follows $M \models^{pa} [\Sigma][\Sigma']\psi$ (Lemma 6 is an equivalence).

- *(Failure of) Left monotonicity*

Consider $\Gamma = \{\neg K_{ap}\}$, $\Sigma = \top$, $\psi = p$ and $\varphi = \neg K_{ap}$. Then, we have that

$$\top \models_{\neg K_{ap}}^d \neg K_{ap},$$

however

$$p, \top \not\models_{\neg K_{ap}}^d \neg K_{ap}. \quad \square$$

Theorem 8 ([25]) *Dynamic consequence satisfies cautious monotonicity, left cut and left monotonicity.*

Proof The proof for cautious monotonicity and left cut is the case $\Gamma = \{\top\}$ of Theorem 7. We prove left monotonicity: $\Sigma \models^d \psi$ implies $\varphi, \Sigma \models^d \psi$. But this is obvious, as it is an instantiation of necessitation for the announcement operator. The principle of necessitation is: $\models^{pa} \sigma$ implies $\models^{pa} [\varphi]\sigma$. Therefore, $\models^{pa} [\Sigma]\psi$ implies $\models^{pa} [\varphi][\Sigma]\psi$. \square

2.3 Structural rules for the background conditions

So far we have studied structural rules involving Σ and φ , where Γ is held as a fixed parameter. But we can also consider structural rules where the background conditions are a variable.

A formula φ is *preserved* if, whenever $M, s \models^{pa} \varphi$, then for all $M' \subseteq M$ such that M' contains s , we have that $M', s \models^{pa} \varphi$.

Proposition 9 *Assume that $\Gamma \models^{pa} \varphi$. Then:*

1. $\Sigma \models_{\Gamma}^d \varphi$ may not hold;
2. $\Sigma \models_{\Gamma}^d \varphi$ whenever the formulas in Γ are preserved;
3. $\Sigma \models_{\Gamma}^d \varphi$ whenever the formulas in Γ are boolean.

Proof

1. Let $\Gamma = \{\neg K_a p\}$, and $\Sigma = p$. Assuming that $M \models^{pa} \neg K_a p$, then as long as there is a p -state in M , we have $M \not\models^{pa} [p]\neg K_a p$. Therefore $p \not\models_{\neg K_a p}^d \neg K_a p$.
2. If formulas are preserved under submodels, then the successive announcements of Σ preserve the truth of all formulas in Γ : $\Sigma \models_{\Gamma}^d \Gamma$. From that and $\Gamma \models^{pa} \varphi$ we get $\Sigma \models_{\Gamma}^d \varphi$.
3. If formulas of Γ are boolean, then they do not change their truth value by successive updates. In other words, they are preserved. So again, any of their \models^{pa} -consequences remains the case. □

Proposition 10 *Let $\Gamma, \Delta \subseteq \mathcal{L}^{el}$.*

- *Let $\Delta \models^{pa} \Gamma$. Then $\Sigma \models_{\Gamma}^d \varphi$ implies $\Sigma \models_{\Delta}^d \varphi$.*
- *Let $\Gamma \subseteq \Delta$. Then $\Sigma \models_{\Gamma}^d \varphi$ implies $\Sigma \models_{\Delta}^d \varphi$. *background monotonicity**

Proof We prove the first. Suppose $\Sigma \models_{\Gamma}^d \varphi$. Let a model M be such that $M \models^{pa} \Delta$. From that and $\Delta \models^{pa} \Gamma$ follows $M \models^{pa} \Gamma$, so $M \models^{pa} [\Sigma]\varphi$. Therefore $\Sigma \models_{\Delta}^d \varphi$. The second item is a special case of the first item. □

One can also consider structural rules that involve variation in the premises and the background conditions simultaneously.

Proposition 11 *The following structural rules hold for conditional dynamic consequence.*

- *Let $\Gamma \models^{pa} \psi$. Then $\psi, \Sigma \models_{\Gamma}^d \varphi$ implies $\Sigma \models_{\Gamma}^d \varphi$. *cautious contraction**
- *$\models_{\Sigma}^d \varphi$ iff $\Sigma \models_{\Sigma}^d \varphi$.*

Proof

- Assume $\psi, \Sigma \models_{\Gamma}^d \varphi$. Let M be given such that $M \models^{pa} \Gamma$. Then, by definition of conditional dynamic consequence, $M \models^{pa} [\psi][\Sigma]\varphi$. Further, from $M \models^{pa} \Gamma$ and $\Gamma \models^{pa} \psi$ follows $M \models^{pa} \psi$, so that announcing ψ is uninformative in M . Therefore, also $M \models^{pa} [\Sigma]\varphi$. We obtained the required $\Sigma \models_{\Gamma}^d \varphi$.
- The statement follows directly from the observation that, for every model M such that $M \models^{pa} \Sigma$, $M = M|\Sigma$. □

When taking the *sequence* Σ as a *set* of background conditions, as in Proposition 11, we abstract from the order of the formulas in that sequence. But on condition that the formulas in Σ are model validities, their order in the sequence no longer matters. For example, let $\Sigma = \{\varphi_1, \varphi_2\}$, then we have for any M with $M \models \Sigma$ (i.e., $M \models \varphi_1$ and $M \models \varphi_2$) that $M|\varphi_1|\varphi_2 = M|\varphi_2|\varphi_1 = M$. We therefore even have that $\pi(\Sigma) \models_{\Sigma}^d \varphi$ iff $\models_{\Sigma}^d \varphi$, where π is any permutation of Σ .

2.4 Announcing background conditions

In Proposition 11 we investigated what happens when the background conditions of dynamic consequence are announced, thus constituting the premises of that dynamic consequence. But we kept them as background conditions. We merely had an announced sequence that was identical to the background conditions. Are there circumstances under which we can get rid of the background conditions by way of announcing them? This would be similar to the proceedings in [18], one of our motivating publications: there, the background assumptions can be added to the premise set. Such a procedure would be of interest, because conditional dynamic consequence would then be reduced to unconditional dynamic consequence. In this section we investigate when this can be done.

Example 12 *Example 5 demonstrated that we cannot always make the background condition into a first announcement: $p \models_{p \wedge \neg Kp}^d \neg p$, whereas $p \wedge \neg Kp, p \not\models^d \neg p$. But sometimes we can. For example, $\models_p^d Kp$ but also $p \models^d Kp$. Both are obvious. For the first we observe that any model on which p is valid also validates Kp . The second amounts to the validity $\models^{pa} [p]Kp$.*

In order to investigate when the background conditions can be announced instead, we have to assume that there are only finitely many of them, because the sequence of the premises of dynamic consequence is required to be finite. We can replace a finite set of conditions Γ by a single condition, their conjunction $\gamma ::= \bigwedge_{\psi \in \Gamma} \psi$. In the remainder of this section we assume a background consisting of one condition γ .

The reason that we cannot put the condition $\gamma = p \wedge \neg Kp$ in Example 5 ahead of the announced premises, seems to be that it cannot be a model validity. We call a formula *globally satisfiable* if there is a model M such that $M \models^{pa} \gamma$. Globally satisfiable is not the

same as satisfiable (there is a model M and a state $s \in M$ such that $M, s \models^{pa} \gamma$). Formula $p \wedge \neg Kp$ is satisfiable but not globally satisfiable.

The following counterexample shows that even if a background condition is globally satisfiable, we may not always announce it first, instead. This is becoming less trivial.

Example 13 Consider the muddy children puzzle, and the formalization of the event ‘nobody steps forward’. Let us take the case of three children. The formula is $\neg(K_a m_a \vee K_a \neg m_a) \wedge \neg(K_b m_b \vee K_b \neg m_b) \wedge \neg(K_c m_c \vee K_c \neg m_c)$, where propositional variable m_i stands for ‘child i is muddy’, for $i = a, b, c$. We abbreviate this formula as ψ_{stand} . A typical case is when a and b are muddy, such that after father’s initial announcement ‘there is at least one muddy child’ ($m_a \vee m_b \vee m_c$), and after his subsequent request that those who know whether they are muddy step forward, nobody steps forward. Then, after he repeats his request, the muddy children a and b will step forward. Let M be the initial model for the muddy children problem, wherein every child i is only uncertain about the value of m_i , and let 110 be the state where a and b are muddy and c not. Consider the model $M' = M|(m_1 \vee m_2 \vee m_3)$. We then have $M'|\psi_{stand}, 110 \not\models^{pa} \psi_{stand}$, and therefore

$$\psi_{stand} \not\models^d \psi_{stand}.$$

On the other hand, if ψ_{stand} is already valid on a model, then it remains valid, so, trivially:

$$\models_{\psi_{stand}}^d \psi_{stand}.$$

The formula ψ_{stand} is globally satisfiable, for example on the model consisting of all valuations of the three variables m_a, m_b, m_c , and universal relations for all three agents a, b, c (so-called ‘blissfull ignorance’: common knowledge of maximal factual ignorance for all agents).

The result of Example 13 can be generalized. A formula φ is *successful* iff $\models^{pa} [\varphi]\varphi$ (see [26], based on [14]). (We note that all contradictions are successful.) A formula is *unsuccessful* if it is not successful. If a formula is unsuccessful, there is a model M and a state $s \in M$ such that φ becomes false after being announced: $M, s \models^{pa} \varphi$, but $M|\varphi, s \not\models^{pa} \varphi$. (We note that all unsuccessful formulas are satisfiable.) Formula ψ_{stand} in Example 13 is unsuccessful, because after nobody steps forward (when it is true), a and b will step forward (it is false).

Proposition 14 Let $\varphi \in \mathcal{L}^{el}$ be unsuccessful and globally satisfiable. Then $\varphi \not\models^d \varphi$ but $\models_{\varphi}^d \varphi$.

Proof Both are obvious. To demonstrate $\varphi \not\models^d \varphi$, take a model M' as in Example 13. This model must exist, because φ is unsuccessful. Whereas $\models_{\varphi}^d \varphi$ has a tautological reading by definition: ‘for all M , $M \models^{pa} \varphi$ implies $M \models^{pa} \varphi$ ’. \square

Theorem 15 Let $\gamma \in \mathcal{L}^{el}$ be successful and satisfiable. Then $\Sigma \models_{\gamma}^d \varphi$ iff $\gamma, \Sigma \models^d \varphi$.

Proof From left to right: The satisfiable successful formulas are also globally satisfiable. Namely, take any M , then φ is a model validity on a non-empty $M|\varphi$. (And at least one such M must exist because φ is satisfiable.) This also answers the proof obligation of the non-trivial direction from right to left. However, let us be explicit.

From right to left: We assume the right-hand side $\gamma, \Sigma \models^d \varphi$. Now let $M \models^{pa} \gamma$, and let $s \in M$. We then have, from the assumption, that $M, s \models^{pa} [\gamma][\Sigma]\varphi$. As we also have $M, s \models^{pa} \gamma$, we get $M|\gamma, s \models^{pa} [\Sigma]\varphi$. As $M|\gamma = M$, $M, s \models^{pa} [\Sigma]\varphi$. As s was arbitrary, $M \models^{pa} [\Sigma]\varphi$. \square

A few afterthoughts to finish this section. Firstly: can ‘successful’ or ‘globally satisfiable’ be characterized in the logical language?

If a formula φ is valid, then it is also known to be valid, i.e., $K\varphi$ is also valid. *Globally satisfiable* seems to have the same meaning as many interpretations of *knowable* [13], a minefield we would rather not tread in this work—but which also has been given meaning in the context of public announcement logic. There is no syntactic characterization of knowable formulas to our knowledge.

The successful formulas have been characterized for the single-agent case in [16]. The multi-agent case is an open question. It is known that the *positive* formulas, the fragment defined by $\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid K_a\varphi$, are successful [26]. This goes back to [2], who show that the positive formulas characterize the preserved formulas; and it is obvious that preserved formulas are successful.

Secondly, let us not forget the following. This search to get rid of the condition in dynamic consequence may be an interesting technical exercise, but it loses sight of the main reason to consider background conditions, namely that logical dynamics is often performed on a given system that can be described by a set of commonly known background conditions. So that having explicit background conditions in a dynamic consequence relation really implements our modelling intuitions.

Even when such conditions may in principle serve as announcements, this should not so much be seen a way to generate that model by model restriction, but as elimination of non-preferred models that do not satisfy the conditions. As the conditions are commonly known, no actual model restriction is taking place! In that sense, background conditions as initial announcements are rather like PDL test programs, but then for model validities.

Example 16 *In the setting of card deals, consider the axiom describing that every player knows her own card (as in Example 4), with clauses like $0_a \rightarrow K_a 0_a$ (if a holds card 0, she knows it). This knowledge results from that player picking up her cards after they have been dealt, in some model M . This is a semi-public action, not an announcement, resulting in a refinement of the accessibility relation for a . We can instead take this axiom as an announcement (it is a positive formula, read it as $\neg 0_a \vee K_a 0_a$). Then it means that only those models ‘survive’ the announcement that already globally satisfy the axiom. The model M above prior to picking up those cards is not among those, because it does not satisfy $0_a \rightarrow K_a 0_a$ (yet).*

We now continue this exposition with a number of extensions and variations on the proposed framework for dynamic consequence: what happens when we add common knowledge, what if we allow infinite sequences of premises, and what about other forms of dynamics, in which case we look at *soft* public announcements.

3 Common knowledge

We proposed conditional dynamic consequence as a relation between a finite set of epistemic formulas and an epistemic formula, given background conditions of epistemic formulas. If we allow common knowledge formulas as well, conditional dynamic consequence is definable as dynamic consequence (subject to reasonable constraints). We list some basic results about common knowledge after which we demonstrate the definability.

A model M is *connected* if for any $s, s' \in M$ we have that $s \sim_A s'$. If M is connected, and $s \in M$, then $M \models^{pa} \varphi$ iff $M, s \models^{pa} C\varphi$. If a proposition is common knowledge, then announcing it does not make an informative difference: if $M, s \models^{pa} C\varphi$, then the connected part of M containing s is the connected part of $M|C\varphi$ containing s .⁵

Proposition 17

1. $\Sigma \models_{\Gamma}^d \varphi$ iff $\Sigma \models_{\Gamma}^d C\varphi$;
2. $\Sigma \models_{\Gamma}^d \varphi$ iff $\Sigma \models_{C\Gamma}^d \varphi$;
3. $\Sigma \models_{\Gamma}^d \varphi$ implies $C\Sigma \models_{\Gamma}^d \varphi$;
4. $C\Sigma \models_{\Gamma}^d \varphi$ does not imply $\Sigma \models_{\Gamma}^d \varphi$.

Proof

1. The non-trivial direction is from left to right. Let $M \models^{pa} [\Sigma]\varphi$. Suppose the sequence Σ can be executed somewhere in M and consider the model restriction $M|\Sigma$. Then $M|\Sigma \models^{pa} \varphi$. This is equivalent to $M|\Sigma \models^{pa} C\varphi$. Therefore, $M \models^{pa} [\Sigma]C\varphi$.
2. Here we use that on any model M , $M \models^{pa} \Gamma$ iff $M \models^{pa} C\Gamma$.
3. Assume a model M such that $M \models^{pa} \Gamma$. We then have by definition of conditional dynamic consequence that $M \models^{pa} [\Sigma]\varphi$. We prove that $M \models^{pa} [C\Sigma]\varphi$ by induction on the length of Σ . Let $\Sigma = \psi, \Sigma'$ (the case $\Sigma = \emptyset$ is trivial). We then have to prove that $M \models^{pa} [C\psi][C\Sigma']\varphi$. Assume a $s \in M$ such that $M, s \models^{pa} C\psi$. Then $C\psi$ does not result in a model restriction on the connected part of M containing s , and ψ

⁵We really want to write $M = M|C\varphi$ here, but if M is not connected this may not hold. We then still have that $(M, s) \Leftrightarrow (M|C\varphi, s)$, where \Leftrightarrow stands for ‘bisimilar’, a notion of sameness of structures implying that they cannot be distinguished in the logical language. Alternatively, if M is not a connected model then similar observations as for common knowledge C can be made for the universal modality U defined as: $M, s \models^{pa} U\varphi$ iff for all $s \in M$ it holds that $M, s \models^{pa} \varphi$.

must already have been valid on that part. (I.e.: $M|C\psi, s$ and $M|\psi, s$ and M, s are the same on the connected part of the model that contains s .) As from $M, s \models^{pa} C\psi$ it also follows that $M, s \models^{pa} \psi$, from that and the assumption $M, s \models^{pa} [\psi][\Sigma']\varphi$ it follows that $M|\psi, s \models^{pa} [\Sigma']\varphi$. From the latter it follows by inductive hypothesis that $M|\psi, s \models^{pa} [C\Sigma']\varphi$. Finally, in view of the already established identification of $M|C\psi, s$ with $M|\psi, s$, it then follows that $M|C\psi, s \models^{pa} [C\Sigma']\varphi$, as required.

4. For example, we have that $C(p \wedge \neg Kp) \models^d \neg p$. This is for the trivial reason that $C(p \wedge \neg Kp)$ is inconsistent. But, of course, $p \wedge \neg Kp \not\models^d \neg p$. \square

Theorem 18 *With respect to the language \mathcal{L}_C^{el} of epistemic logic with common knowledge, dynamic consequence satisfies cautious monotonicity, left cut and left monotonicity, and conditional dynamic consequence satisfies cautious monotonicity and left cut, but not left monotonicity.*

Proof The proofs of the structural properties of Theorems 7 and 8 only use the announcement modality of the logical language, and not other modalities. The counterexample establishing failure on left monotonicity for conditional dynamic consequence (that uses the knowledge modality) is still a counterexample in the presence of common knowledge. \square

As the following theorem is about announcing background conditions, we have to assume that this is a finite set. It therefore can be represented by a single condition γ .

Theorem 19 *In \mathcal{L}_C^{pa} , on the class of connected models, conditional dynamic consequence is definable as unconditional dynamic consequence:*

$$\Sigma \models_{\gamma}^d \varphi \quad \text{iff} \quad C\gamma, \Sigma \models^d \varphi.$$

Proof ‘ \Rightarrow ’: Let M and $s \in M$ be given, and assume that $M, s \models^{pa} C\gamma$. Then, because M is connected, $M \models \gamma$, so, using the assumption, we conclude $M \models^{pa} [\Sigma]\varphi$. In particular, we get $M, s \models^{pa} [\Sigma]\varphi$, our remaining proof obligation.

‘ \Leftarrow ’: Let M be given and let $M \models^{pa} \gamma$. We then also have $M \models^{pa} C\gamma$. So, for any $s \in M$: $M, s \models^{pa} C\gamma$. Now using the assumption, we conclude $M|C\gamma, s \models^{pa} [\Sigma]\varphi$. As this was for arbitrary s , and as $M|C\gamma = M$, we get the required $M \models^{pa} [\Sigma]\varphi$. \square

4 Infinite sequences of premises

A natural generalization of dynamic consequence $\Sigma \models^d \varphi$ is to infinite sequences of premises Σ ; and a further natural generalization of conditional dynamic consequence $\Sigma \models_{\Gamma}^d \varphi$ would then be to investigate when infinite sets of background conditions Γ can be announced instead. As $[\Sigma]\varphi$ is undefined for infinite sequences we need to define $\Sigma \models^d \varphi$ in an alternative (obvious) way.

For an infinite sequence of premises $\Sigma = \psi_1, \psi_2, \dots$ we define the initial subsequence Σ_n of length n as follows: $\Sigma_1 = \psi_1$, $\Sigma_{n+1} = \Sigma_n, \psi_{n+1}$.

Definition 20 (Dynamic consequence with infinite premises) Let Σ be an infinite sequence of premises. Then

$$\Sigma \models^d \varphi \quad \text{iff} \quad \text{for all models } M : M|\Sigma \models^{pa} \varphi$$

where $M|\Sigma$ is defined (if non-empty) as the restriction of the model M to the domain consisting of the states $s \in M$ such that for all $n \in \mathbb{N}$, $M, s \models^{pa} \langle \Sigma_n \rangle \top$.

Alternatively, and by all means equivalently, we could think of $M|\Sigma$ defined as (a non-empty) $\bigcap_{n \in \mathbb{N}} M_n$ where $M_0 = M$ and $M_n = M|\Sigma_n$.⁶

Definition 21 (Compactness) Dynamic consequence is compact if $\Sigma \models^d \varphi$ implies that there is an n such that $\Sigma_n \models^d \varphi$.

We have few results on this, and an unproven conjecture—namely that it is compact. We thought this of sufficient interest to inform the reader. It is not a ghost hunt, because there are indeed *informative infinite sequences of announcements*, i.e., such that, for any n , there is an $m > n$ with $M|\Sigma_n \neq M|\Sigma_m$.

Example 22 Given a set of variables $P = \{p_1, p_2, \dots\}$, the announcement sequence $\Sigma = p_1, p_2, \dots$ is informative. On the model M where an agent is uncertain between all valuations of these variables, every next announcement will make an informative difference. The agent will know the valuation after the infinite sequence, but not after any initial sequence.

However, given a formula φ , that formula will contain only a finite set $P' \subseteq P$ of propositional variables. Let $p_m \in P'$ be the maximum (given the sequence Σ). It can be easily shown that on a model where the valuation of the variables in a given set is constant (the same in all states), a formula containing only those variables is either valid on the model, or its negation is valid. Therefore, $\Sigma \models^d \varphi$ implies $\Sigma_m \models^d \varphi$.

More complex examples can be made for infinite sequences of announcements containing only finitely many propositional variables, but (stacks of) modalities.

If we widen our scope to more complex events than public announcements, there are ever informative infinite sequences of events that consist of an infinitely often repeated

⁶A problem of this semantics for infinite sequences of announcements is that properties of the model M or of the sequence Σ may be of little help in determining the truth of a given formula φ in $M|\Sigma$. Alternatively and by no means equivalently we could think of $M|\Sigma$ as defined (for convenience of presentation, by way of its semantics) as follows:

$$M|\Sigma \models^{pa} \varphi \quad \text{iff there is an } n \text{ such that for all } m \geq n : M|\Sigma_m \models^{pa} \varphi$$

This leads to very different results! Consider the typical rooted ‘spider’ model with n -length arms/branches (with asymmetric access) for each $n \in \mathbb{N}$, and an announcement removing the leaf of each branch. The result of each such announcement is an isomorphic spider model: the branch of length 1 has now been eliminated, but the branch of length 2 has now length 1. And so on, for every branch. So with the ‘different’ semantics, $M = M|\Sigma$. But after repeating the announcement infinitely often only the singleton root state remains. So with the proposed ‘intersection of all $M|\Sigma_m$ ’-semantics, $M|\Sigma$ is that singleton.

single event only. Intriguing results on this are presented by Sadzik in [21].⁷ Here again, the observation that for some given model M of information there is an infinite informative sequence ψ_1, ψ_2, \dots of announcements, as Sadzik does, is one thing. Whereas the question whether, given a formula φ , there is an infinite informative sequence ψ_1, ψ_2, \dots of announcements such that for some model M the formula φ never becomes a validity, is another thing: it is a weaker statement.⁸

Before we proceed, let us contrast our problem with the known compactness results for these logics. For the consequence relation \models^{pa} the results are that the logic of public announcements with common knowledge is not compact, and that the logic of public announcements without common knowledge is compact [8]. These results expand the results that epistemic logic without common knowledge is compact and epistemic logic with common knowledge is not compact (summarized in [12]). This is the standard set-based notion of compactness: If $\Gamma \models^{pa} \psi$ then there is a finite subset $\Delta \subseteq \Gamma$ such that $\Delta \models^{pa} \psi$. A typical example is the set $\{p, K_a p, K_b p, K_a K_b p, \dots, \neg C p\}$ (any finite stack of two alternating modalities). This set is unsatisfiable but any finite subset is satisfiable.

To use the compactness result for epistemic logic, we need to define a relevant unsatisfiable infinite set. Now consider this. Suppose in order to derive a contradiction that dynamic consequence is not compact. Assume that $\Sigma \models^d \varphi$ but that $\Sigma_n \not\models^d \varphi$ for all n . In other words, for all $n \in \mathbb{N}$, we have that $\not\models^{pa} [\Sigma_n] \varphi$, and therefore:

$$\text{For all } n \in \mathbb{N}, \langle \Sigma_n \rangle \neg \varphi \text{ is satisfiable.} \quad (i)$$

Somehow, we would like to arrive from here at

$$\text{The set } \{ \langle \Sigma_n \rangle \neg \varphi \mid n \in \mathbb{N} \} \text{ is satisfiable.} \quad (ii)$$

But (i) does not imply (ii). Statement (i) says that for all $n \in \mathbb{N}$ we can find a model M_n and a state $s_n \in M_n$ such that: $M_n, s_n \models^{pa} \langle \Sigma_n \rangle \neg \varphi$. Whereas (ii) says that there is a (single) model M and a state $s \in M$ such that for all $n \in \mathbb{N}$: $M, s \models^{pa} \langle \Sigma_n \rangle \neg \varphi$. This pointed model (M, s) does it for the whole set. And even if we get there, and we have an (M, s) satisfying (ii), we cannot know if φ is true or false there: φ might easily be false after every announced Σ_n but still true after Σ .⁹

We only were able to get (ii) from (i) for conclusions of dynamic consequence of a specific form, not for any formula. We now list these relevant results and the conjecture.

Proposition 23 *If $\varphi \in \mathcal{L}^{el}$ with $\neg \varphi$ preserved, then $\Sigma \models^d \varphi$ implies that there is $n \in \mathbb{N}$ such that the following formula is pa-valid:*

$$[\Sigma_1] \varphi \vee [\Sigma_2] \varphi \vee \dots \vee [\Sigma_n] \varphi$$

⁷Probably the most publishable but not formally published result in dynamic epistemic logic.

⁸In [21, Section 4] it is proved that for a given formula φ , events with boolean preconditions, when infinitely repeated, have at some stage no longer informative consequences. The examples given involve non-public events (events that are not public announcements). This result suggests that, in principle, we have a compactness result of the form: If $\psi, \psi, \dots \models^d \varphi$ and ψ is boolean, then there is an n such that $\psi, \psi, \dots, \psi \models^d \varphi$ (where ψ is announced n times). But if ψ is boolean, then for any M , $M \mid \psi \models \psi$, so we would have already had $\psi \models^d \varphi$ straightaway.

⁹See footnote 6.

Proof Towards a contradiction, assume that for each $n \in \mathbb{N}$ the formula

$$[\Sigma_1]\varphi \vee [\Sigma_2]\varphi \vee \dots \vee [\Sigma_n]\varphi$$

is not pa -valid. Then, for each $n \in \mathbb{N}$ the formula

$$\neg([\Sigma_1]\varphi \vee [\Sigma_2]\varphi \vee \dots \vee [\Sigma_n]\varphi)$$

is pa -satisfiable. That is, for each $n \in \mathbb{N}$ the formula

$$(\star) \langle \Sigma_1 \rangle \neg\varphi \wedge \langle \Sigma_2 \rangle \neg\varphi \wedge \dots \wedge \langle \Sigma_n \rangle \neg\varphi$$

is pa -satisfiable. Now consider the set $S = \{\langle \Sigma_n \rangle \neg\varphi \mid n \in \mathbb{N}\}$. Let $S' \subset S$ be any finite subset. The set S' will have a maximal element $\langle \Sigma_m \rangle \neg\varphi$. By (\star) let M and $s \in M$ be such that

$$M, s \models^{pa} \langle \Sigma_1 \rangle \neg\varphi \wedge \langle \Sigma_2 \rangle \neg\varphi \wedge \dots \wedge \langle \Sigma_m \rangle \neg\varphi.$$

Therefore M, s satisfies the entire set $S'' = \{\langle \Sigma_n \rangle \neg\varphi \mid n \leq m\}$. As $S' \subseteq S''$, S' is also satisfiable (by (M, s)).

So, any finite subset of S is satisfiable. From the compactness of public announcement logic without common knowledge it follows that S is satisfiable as well. Thus, there exist N and $t \in N$ satisfying that for all $n \in \mathbb{N}$:

1. $N, t \models^{pa} \langle \Sigma_n \rangle \top$, and
2. $N, t \models^{pa} \langle \Sigma_n \rangle \neg\varphi$.

From the first item follows that $t \in N|\Sigma$. From the second item follows that $N|\Sigma, t \models^{pa} \neg\varphi$. For that, observe that $N|\Sigma_1, t \models^{pa} \neg\varphi$ and $\neg\varphi$ is preserved. But this contradicts $\Sigma \models^d \varphi$. \square

Proposition 23 gives us almost *but not quite* what we want. Although we can get a n such that $\models^{pa} [\Sigma_1]\varphi \vee [\Sigma_2]\varphi \vee \dots \vee [\Sigma_n]\varphi$, for compactness (given the restrictions on φ) we would have liked $\Sigma_n \models^d \varphi$, i.e., $\models^{pa} [\Sigma_n]\varphi$. We cannot get that from the former disjunction, unless even more restrictions are satisfied:

Corollary 24 *If φ is boolean, then $\Sigma \models^d \varphi$ implies that there is an n such that $\Sigma_n \models^d \varphi$.*

Proof The result follows from Proposition 23 together with the following facts.

1. All booleans and their negations are positive formulas (see Section 2.4), and therefore preserved.
2. If φ is preserved, the following formula is pa -valid

$$([\Sigma_1]\varphi \vee [\Sigma_2]\varphi \vee \dots \vee [\Sigma_n]\varphi) \rightarrow [\Sigma_n]\varphi$$

To see this, let M and $s \in M$ be such that $M, s \models^{pa} [\Sigma_1]\varphi \vee [\Sigma_2]\varphi \vee \dots \vee [\Sigma_n]\varphi$. Then there is $i \leq n$ such that $M, s \models^{pa} [\Sigma_i]\varphi$. Now as φ is preserved, we also have that $M, s \models^{pa} [\Sigma_n]\varphi$. \square

Conjecture 25 *Dynamic consequence is compact.*

5 Dynamic consequence for soft public announcements

An alternative to the ‘hard’ public announcements that we have been dealing with so far are the so-called ‘soft’ public announcements. In a soft public announcement of φ , the agents are willing to incorporate φ , but only tentatively so. When they are later informed of something that would be in contradiction with (the informative consequences of) φ , their beliefs will not become inconsistent, but they will change again in some consistency preserving way. A typical example is to compare the hard public announcement with the soft public announcement of a primitive proposition p . Just a hard public announcement of p will make the agent (or agents) know p , a soft public announcement of p will make the agent believe p . After a hard public announcement of p the agent cannot incorporate hard public announcement of $\neg p$: that is in contradiction with his (infallible) knowledge of p . But after a soft public announcement of p the agent can very well incorporate soft public announcement of $\neg p$. After that, he will no longer believe p , but believe $\neg p$. Hard public announcement is a form of belief expansion, and soft public announcement is a form of belief revision.

The dynamic epistemic logic way to implement belief revision as soft public announcement is presented in [24]. For this we need to enrich the epistemic models with plausibility relations that tell us, given two states that an agent considers possible, which of these is more plausible than the other. An agent believes a proposition φ iff φ is true in the most plausible states. A way to implement this (following van Benthem’s presentation) is to have, apart from an equivalence relation \sim_a for every agent a , also a plausibility relation $<_a$ for every agent a (that is a well-preorder, i.e., a well-founded, reflexive, and transitive relation), such that we can define $\prec_a ::= <_a \cap \sim_a$ (\leq_a is the reflexive closure of $<_a$, and \preceq_a is the reflexive closure of \prec_a). This then defines the plausibility order per equivalence class. The well-preorder requirement is to ensure that there are most plausible states in every class (most plausible is $<_a$ -minimal), so that belief is always consistent. Such models $(S, \sim, <, V)$ are often called *plausibility epistemic models*.

The simplest language to be considered under this setting is

$$\mathcal{L}^{bpa} \quad \ni \quad \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid B_a\varphi \mid \llbracket\varphi\rrbracket\varphi$$

where $B_a\varphi$ stands for ‘agent a believes φ ’ and $\llbracket\varphi\rrbracket\psi$ stands for ‘after soft public announcement of φ agent a believes that ψ ’. The language without soft public announcement operators is \mathcal{L}^{bl} (the language of doxastic logic). There are various extensions of the language \mathcal{L}^{bpa} : with knowledge, with conditional belief—essential for a complete axiomatization—, with group epistemic operators or group epistemic operators, ... We need not consider these extensions here. The semantics of the two novel modalities is now defined as follows.

$$\begin{aligned} M, s \models^{bpa} B_a\varphi & \text{ iff for every } t \text{ such that } s \sim_a t \text{ and } t \leq_a u \text{ for all } u \sim_a s, M, t \models^{bpa} \varphi \\ M, s \models^{bpa} \llbracket\varphi\rrbracket\psi & \text{ iff } M^\varphi, s \models^{bpa} \psi \end{aligned}$$

where M^φ is as M except that the plausibility relation $<'$ is defined as

$$\begin{aligned}
t <'_a t' \quad \text{iff} \quad & t <_a t', \text{ and } M, t \models^{bpa} \varphi \text{ and } M, t' \models^{bpa} \varphi, \\
& \text{or} \\
& M, t \models^{bpa} \varphi \text{ and } M, t' \not\models^{bpa} \varphi, \\
& \text{or} \\
& t <_a t', \text{ and } M, t \not\models^{bpa} \varphi \text{ and } M, t' \not\models^{bpa} \varphi.
\end{aligned}$$

The idea behind this operation is that after soft public announcement of φ , all φ -states are considered more plausible than all $\neg\varphi$ -states, and that the existing plausibility relation among the φ -states and also among the $\neg\varphi$ -states is otherwise preserved.

Example 26 *Consider a controller room with a light and a fan, the state of the light is represented by a proposition p (light on/off) and the state of the fan by a proposition q . A model representing the uncertainty of the controller contains all four valuations pq , $p\neg q$, $\neg pq$ and $\neg p\neg q$. The controller's beliefs are as follows: $pq <_a \neg pq <_a p\neg q =_a \neg p\neg q$; in other words: he considers it most plausible that the light and fan are both functional, slightly less so that at least the fan is working (but the light off), and even less so that the fan is not working (where he is indifferent about the state of the light). The controller believes that the light and fan are both functional: $B_a p \wedge B_a q$. To ascertain the state of affairs, he can make a tentative observation namely that the light is off. It is tentative, because it is a bit hard to see, as there is daylight outside the controller room. As a result of that, all $\neg p$ states now become more plausible than all p states, but the order within the $\neg p$ states and the order within the p states is preserved. Therefore, the new plausibility relation is: $\neg pq <_a \neg p\neg q <_a pq <_a p\neg q$. The controller now believes that the light is off but the fan still on: $B_a \neg p \wedge B_a q$. Before the soft public announcement (i.e., tentative observation) of $\neg p$, he believed that p , but now, afterwards, he believes that $\neg p$. So, referring to the initial state of affairs, we have that $B_a p \wedge \llbracket \neg p \rrbracket B_a \neg p$ is true there.*

In [25], van Benthem also proposes dynamic consequence for soft public announcement, namely as $\Sigma \models \varphi$ iff $\models^{bpa} \llbracket \Sigma \rrbracket B\varphi$ (after soft public announcement of Σ , φ is always believed) and as $\Sigma \models \varphi$ iff $\models^{bpa} [\Sigma] B\varphi$ (after hard public announcement of Σ , φ is always believed); where in the second case we further need to extend the language \mathcal{L}^{bpa} with a clause for hard public announcement and a corresponding clause in the semantics. A peculiarity of the former is that even for atomic propositions belief does not persist, as we now have $p \models p$ but $p, \neg p \not\models p$. This made him predict (but not investigate) different structural properties of such dynamic consequence relations.

We followed a different route than van Benthem. Although $\models^{pa} \varphi$ is equivalent to $\models^{pa} K\varphi$, clearly $\models^{bpa} \varphi$ is not equivalent to $\models^{bpa} B\varphi$. If one instead requires model validity of φ , not $B\varphi$, the obvious definition is as follows. We give this straightaway in unconditional and conditional version, and for multiple agents.

Definition 27 (Soft dynamic consequence) *Let $\varphi \in \mathcal{L}^{bl}$ and $\Sigma \subseteq \mathcal{L}^{bl}$.*

$$\begin{aligned} \Sigma \models^{bd} \varphi \quad \text{iff} \quad & \models^{bpa} \llbracket \Sigma \rrbracket \varphi \\ & \text{soft (unconditional) dynamic consequence} \\ \Sigma \models_{\Gamma}^{bd} \varphi \quad \text{iff} \quad & \text{for all } M : M \models^{bpa} \Gamma \text{ implies } M \models^{bpa} \llbracket \Sigma \rrbracket \varphi \\ & \text{soft conditional dynamic consequence} \end{aligned}$$

The conditional version is now of more interest than before, as the conditions are *knowledge* conditions (commonly known background conditions) whereas the premises of the dynamic consequence relation are soft public announcements. Soft public announcements affect beliefs but not (factual) knowledge.

Surprisingly, for soft dynamic consequence we obtain the same results as for (hard) dynamic consequence. In the first place, the classical substructural rules also fail for soft dynamic consequence. This may be not too surprising. But we considered it surprising that they fail on the same counterexamples as for hard dynamic consequence—we merely need to replace knowledge by belief.

- Reflexivity fails: $\llbracket p \wedge \neg B_a p \rrbracket \neg(p \wedge \neg B_a p)$ is valid, so $p \wedge \neg B_a p \not\models^{bd} p \wedge \neg B_a p$.
- Contraction fails: $p \wedge \neg B_a p, p \wedge \neg B_a p \models^{bd} \perp$, but $p \wedge \neg B_a p \not\models^{bd} \perp$.
- Permutation fails: $p, \neg B_a p \models^{bd} \perp$, but $\neg B_a p, p \not\models^{bd} \perp$.
- Monotonicity fails: $\neg B_a p \models^{bd} \neg B_a p$, however, $\neg B_a p, p \not\models^{bd} \neg B_a p$.
- Cut fails. We now have that $\neg p \models^{bd} B_a \neg p$, and that $\neg B_a \neg p, B_a \neg p \models^{bd} \perp$; but $\neg B_a \neg p, \neg p \not\models^{bd} \perp$.

Secondly, and more surprisingly so, soft dynamic consequence satisfies (at least) the same structural properties as hard dynamic consequence and conditional soft dynamic consequence satisfies (at least) the same structural properties as conditional (hard) dynamic consequence. As before (see Lemma 6), a helpful lemma is:

Lemma 28 *If $M \models^{bpa} \llbracket \Sigma \rrbracket \varphi$, then for all χ , $M \models^{bpa} \llbracket \Sigma \rrbracket \llbracket \varphi \rrbracket \chi \leftrightarrow \llbracket \Sigma \rrbracket \chi$.*

Proof This is now true because soft public announcement of a model validity (also) has no informative consequence. Of the three cases of the computation of the new plausibility relation, only the first applies: on the part of the model that satisfies the soft public announcement formula, the plausibility relation remains the same. That is, for a formula that already is a model validity: it remains the same on the entire model. \square

Theorem 29 *Soft dynamic consequence satisfies cautious monotonicity, left cut and left monotonicity.*

Proof The structural rules cautious monotonicity and left monotonicity are the special case of $\Gamma = \{\top\}$ of the proof for conditional soft dynamic consequence, next.

We prove left monotonicity: $\Sigma \models^{bd} \psi$ implies $\varphi, \Sigma \models^{bd} \psi$. Just as in the proof of Theorem 8, this is an instantiation of necessitation, which also holds for the soft public announcement operator. (Necessitation is: $\models^{bpa} \sigma$ implies $\models^{bpa} \llbracket \varphi \rrbracket \sigma$. Then, take $\sigma = \llbracket \Sigma \rrbracket \psi$.) \square

Theorem 30 *Conditional soft dynamic consequence satisfies cautious monotonicity and left cut but not left monotonicity.*

Proof

- *Cautious monotonicity*

Assume that $\Sigma \models_{\Gamma}^{bd} \varphi$ (i), that $\Sigma, \Sigma' \models_{\Gamma}^{bd} \psi$ (ii), and that, for some model M , $M \models^{bpa} \Gamma$. From the last and assumption (i) follows $M \models^{bpa} \llbracket \Sigma \rrbracket \varphi$. From $M \models^{bpa} \Gamma$ and assumption (ii) follows $M \models^{bpa} \llbracket \Sigma \rrbracket \llbracket \Sigma' \rrbracket \psi$. From $M \models^{bpa} \llbracket \Sigma \rrbracket \varphi$, $M \models^{bpa} \llbracket \Sigma \rrbracket \llbracket \Sigma' \rrbracket \psi$ and Lemma 28, by taking $\chi = \llbracket \Sigma' \rrbracket \psi$, follows $M \models^{bpa} \llbracket \Sigma \rrbracket \llbracket \varphi \rrbracket \llbracket \Sigma' \rrbracket \psi$.

- *Left cut*

Assume $M \models^{bpa} \llbracket \Sigma \rrbracket \varphi$. From $M \models^{bpa} \llbracket \Sigma \rrbracket \llbracket \varphi \rrbracket \llbracket \Sigma' \rrbracket \psi$ also follows $M \models \llbracket \Sigma \rrbracket \llbracket \Sigma' \rrbracket \psi$, as Lemma 28 is an equivalence.

- *(Failure of) Left monotonicity*

The same counterexample as in Theorem 7 holds (replacing K by B). Consider $\Gamma = \{\neg B_a p\}$, $\Sigma = \top$, $\psi = p$ and $\varphi = \neg B_a p$. Then, we have that

$$\top \models_{\neg B_a p}^{bd} \neg B_a p,$$

however

$$p, \top \not\models_{\neg B_a p}^{bd} \neg B_a p. \quad \square$$

We did not investigate if these three rules *represent* this (unconditional) soft dynamic consequence relation in the technical sense of lifting the discussion to abstract transition models, just as in [25] for (hard) dynamic consequence. We can merely observe that we did not find structural rules, except those explicitly formulated in terms of K or B operators, that are satisfied by hard dynamic consequence but not by soft dynamic consequence. (And we can observe that the proofs of the following theorems are all similar: Theorem 7 and 8, Theorem 18, and Theorem 29 and 30.)

In the setting of soft dynamic consequence there is little point in investigating background conditions as announcements. These are now very different phenomena. Background conditions are hard constraints, corresponding, if one wishes, to hard public announcements. Whereas soft dynamic consequence is relative to soft public announcements, a clearly different form of belief revision. We see this as an advantage. In conditional soft dynamic consequence the hard information part is in the background conditions, whereas the soft information part is in the premises. As in many information processing scenarios we have a clear distinction between hard and soft information. Conditional soft dynamic consequence can therefore be seen as an implementation of that role distinction.

6 Conclusions, related work, and future work

Given formulas in multi-agent epistemic logic, we investigated the structural properties of a conditional dynamic consequence relation $\Sigma \models_{\Gamma}^d \varphi$, that stands for ‘on all models that satisfy the set of background conditions Γ , after announcing the sequence Σ , the formula φ always holds’. This consequence relation does not satisfy the so-called classical structural rules reflexivity, contraction, permutation, monotonicity, and cut; it satisfies cautious monotonicity and left cut. In the presence of common knowledge, conditional dynamic consequence is definable from (unconditional) dynamic consequence. We leave for further research whether dynamic consequence is compact, although we prove a ‘disjunctive’ form of compactness for the formulas such that their negation is preserved under submodel restrictions. Our results for (hard) conditional dynamic consequence also hold for soft dynamic consequence. As soft announcement corresponds to radical upgrade in the belief revision community, further research might reveal how conditional dynamic consequence behaves on other update policies, such as conservative upgrade. Our hope is that dynamic consequence relations provide a novel and productive perspective on reasoning about protocols in multi-agent systems.

A tableau calculus for dynamic consequence seems an interesting challenge, where, possibly, the reported calculi for public announcement logic such as [7] may provide inspiration. Recalling the dynamic consequence pattern $\Sigma \models_{\Gamma}^d \varphi$, we focussed on $\Sigma \models_{\Gamma}^{d?} \varphi$: what could be obtained given conditions Γ and a sequence of announcements Σ . But one could equally think of this as $? \models_{\Gamma}^d \varphi$: given background conditions and a desirable goal, which sequence of announcements will always realize that goal? We can either see this as a problem of abduction, as in [1], in a philosophical logical setting. But one can equally well see this as a classical AI planning problem, but in a dynamic epistemic logic, as in [9, 4].

In works such as [3, 4] Aucher and collaborators investigate a consequence relation $\gamma, \psi^d \models \varphi$ where γ is a formula in a logical language (e.g. \mathcal{L}^{el}) that is true in a pointed Kripke structure, ψ^d is a formula in another type of logical language, for actions (the upper index d is ours, not Aucher’s; it is to emphasize its dynamic role), and φ is a formula true in the Kripke structure resulting from executing an action satisfying ψ^d in the Kripke model satisfying γ . For actions that are public announcements we can almost, but not quite, think of ψ^d as the announced formula. This two-typed consequence relation (γ, φ of ‘static’ type, and ψ^d of ‘dynamic’ type) has therefore aspects of dynamic consequence. In the cited works sequent calculi are proposed for this consequence relation. In [6] the approach is generalized to a consequence relation of form $\gamma_1, \psi_1^d, \dots, \gamma_n, \psi_n^d \models \varphi$, i.e., where pairs of static/dynamic premises γ_i, ψ_i^d formalize the successive execution of a sequence of actions satisfying action formulas ψ_i^d , each preceded by a static requirement γ_i . In the forthcoming [5] van Benthem’s dynamic consequence $\psi_1, \dots, \psi_n \models^d \varphi$ is then embedded in this framework, as some such $\gamma_1, \psi_1^d, \dots, \gamma_n, \psi_n^d \models \varphi$. It does not simply replace the announced ψ_i by action formulas ψ_i^d , for example the embedding also employs common knowledge operators (infinitary modalities), but the relation is close enough to use suggestive corresponding notation. In future research we wish interpret our results in Aucher’s

framework.

Acknowledgements

In an early stage where ideas for dynamic consequence floated around, Hans van Ditmarsch acknowledges various discussions with Ramanujam at NIAS Wassenaar in 2007 and with Ramanujam and S.P. Suresh at IMSc in Chennai from 2008 onward. We thank Johan van Benthem for his input. We further thank various participants of the ever evolving LPIS seminar in Seville, from 2009 onward: Laura Leonides, Cristina Barés, David Fernández, Emililio Gómez-Caminero, Ignacio Hernández, Joost Joosten, and Fernando Soler. A previous version of this work appeared in workshop proceedings [10]. Subsequently, we thank Guillaume Aucher, Barteld Kooi, Olivier Roy, and Sonja Smets for further input, at M4M Osuna, in Munich, and in Groningen. Hans van Ditmarsch is also affiliated to IMSc (Institute of Mathematical Sciences), Chennai, as research associate. This work was mainly carried out while Hans van Ditmarsch worked in Seville, supported by project P08-HUM-04159. We acknowledge support from European Research Council grant EPS 313360. We thank the reviewers of the *Review of Symbolic Logic* for their comments.

References

- [1] A. Aliseda. *Abductive Reasoning. Logical Investigations into Discovery and Explanation*. Synthese Library, volume 220. Springer, 2006.
- [2] H. Andréka, I. Németi, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
- [3] G. Aucher. DEL-sequents for progression. *Journal of Applied Non-Classical Logics*, 21(3-4):289–321, 2011.
- [4] G. Aucher. DEL-sequents for regression and epistemic planning. *Journal of Applied Non-Classical Logics*, 22(4):337–367, 2012.
- [5] G. Aucher. Dynamic epistemic logic as a substructural logic. Manuscript to appear in a forthcoming book volume honouring Johan van Benthem, 2013.
- [6] G. Aucher, B. Maubert, and F. Schwarzentruher. Generalized DEL-sequents. In *Proc. of 13th JELIA*, LNCS 7519, pages 54–66, 2012.
- [7] P. Balbiani, H. van Ditmarsch, A. Herzig, and T. de Lima. Tableaux for public announcement logics. *Journal of Logic and Computation*, 20(1):55–76, 2010.
- [8] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proc. of 7th TARK*, pages 43–56, 1998.

- [9] T. Bolander and M.B. Andersen. Epistemic planning for single and multi-agent systems. *Journal of Applied Non-classical Logics*, 21(1):9–34, 2011.
- [10] A. Cordon-Franco, H. van Ditmarsch, D. Fernández-Duque, E. Gomez-Caminero, and A. Nepomuceno-Fernández. Two mischievous dynamic consequence relations. In *Proceedings of the 2nd ILCLI*, pages 197–208. University of the Basque Country Press, 2010.
- [11] J. Etchemendy. *The Concept of Logical Consequence*. Harvard University Press, 1990.
- [12] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge MA, 1995.
- [13] F.B. Fitch. A logical analysis of some value concepts. *The Journal of Symbolic Logic*, 28(2):135–142, 1963.
- [14] J.D. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999. ILLC Dissertation Series DS-1999-01.
- [15] W. Groeneveld. *Logical investigations into dynamic semantics*. PhD thesis, University of Amsterdam, 1995. ILLC Dissertation Series DS-1995-18.
- [16] W. Holliday and T. Icard. Moorean phenomena in epistemic logic. In L. Beklemishev, V. Goranko, and V. Shehtman, editors, *Advances in Modal Logic 8*, pages 178–199. College Publications, 2010.
- [17] M. Kracht. Dynamic semantics. In F. Hamm and T.E. Zimmermann, editors, *Semantics*, pages 217–241, Hamburg, 2002. Buske Verlag.
- [18] D. Makinson. Bridges between classical and nonmonotonic logic. *Logic Journal of the IGPL*, 11:69–96, 2003.
- [19] J.A. Plaza. Logics of public communications. In *Proc. of the 4th ISMIS*, pages 201–216. Oak Ridge National Laboratory, 1989.
- [20] G. Restall. *An Introduction to Substructural Logics*. Routledge, 2000.
- [21] T. Sadzik. Exploring the iterated update universe. ILLC, Amsterdam. Technical report PP-2006-18, 2006.
- [22] A. Tarski. *On the concept of logical consequence*. Oxford University Press, 1936. Reprinted in A. Tarski, *Logic, Semantics, Metamathematics*, 2nd ed.
- [23] J. van Benthem. *Exploring logical dynamics*. CSLI Publications, 1996.
- [24] J. van Benthem. Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

- [25] J. van Benthem. Logical dynamics meets logical pluralism? *Australasian Journal of Logic*, 6:182–209, 2008.
- [26] H. van Ditmarsch and B. Kooi. The secret of my success. *Synthese*, 151:201–232, 2006.
- [27] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Descriptions of game states. In G. Mints and R. Muskens, editors, *Logic, Games, and Constructive Sets*, pages 43–58, Stanford, 2003. CSLI Publications. CSLI Lecture Notes No. 161.
- [28] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [29] F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.